arXiv:1102.1628v2 [math.MG] 14 Sep 2012

APOLLONIAN CIRCLE PACKINGS OF THE HALF-PLANE

MICHAEL CHING AND JOHN R. DOYLE

Abstract. We consider Apollonian circle packings of a half Euclidean plane. We give necessary and sufficient conditions for two such packings to be related by a Euclidean similarity (that is, by translations, reflections, rotations and dilations) and describe explicitly the group of self-similarities of a given packing. We observe that packings with a non-trivial self-similarity correspond to positive real numbers that are the roots of quadratic polynomials with rational coefficients. This is reflected in a close connection between Apollonian circle packings and continued fractions which allows us to completely classify such packings up to similarity.

AMS 2010 Subject Classifications: 52C26; 11A55

Keywords: Apollonian circle packings; similarity; continued fractions

1. Introduction

A circle packing in $\mathbb{R}^2$ is a set of circles in the plane whose interiors (suitably interpreted) are mutually disjoint. An **Apollonian circle packing** $\mathcal{P}$ has the property that for any three mutually tangent circles in $\mathcal{P}$, the two circles in the plane that are tangent to all three of them also lie in $\mathcal{P}$. Note that our notion of circle includes straight lines where we consider parallel lines to be tangent at infinity.

These types of circle packings have been extensively studied by Graham, et al [3, 4, 5, 6], with a focus on those packings for which all the circles have integer curvatures.

There are four basic shapes that an Apollonian packing may take, and these are illustrated in Figure 1. A **bounded** Apollonian packing (Figure 1(a)) is a packing $\mathcal{P}$ for which a single circle in $\mathcal{P}$ bounds the entire packing. Here the 'interior' of the bounding circle is the *unbounded* component of its complement.

A **half-plane** packing (Figure 1(b)) is an Apollonian packing $\mathcal{P}$ for which at least one of its circles is a straight line. The line partitions the plane into two half-planes: one is 'packed' by $\mathcal{P}$, while the other is the 'interior' of the line. A special type of half-plane packing is a **strip** packing (Figure 1(c)), in which two of the circles are (necessarily parallel) lines and the remaining circles lie in the strip between them.

An **unbounded** packing (Figure 1(d)) is an Apollonian packing which contains no bounding circle and no straight line.

Stereographic projection allows us to relate circle packings in the plane to those on the sphere. The four possible configurations in Figure 1 correspond to projection from (a) an interior point, (b) a point on only one circle, (c) a tangency point, and (d) a point not on or inside any circle, respectively.

In this paper, we consider the similarity relation on Apollonian circle packings. A **similarity** is a transformation of the Euclidean plane that preserves ratios of lengths. Such a transformation is

---

Research for this paper was supported by NSF VIGRE grant DMS-0738586.

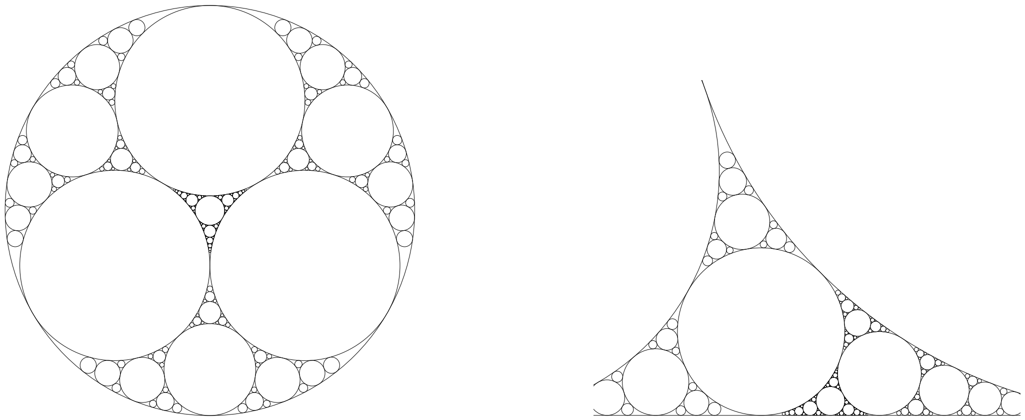

(a) Bounded (b) Half-plane

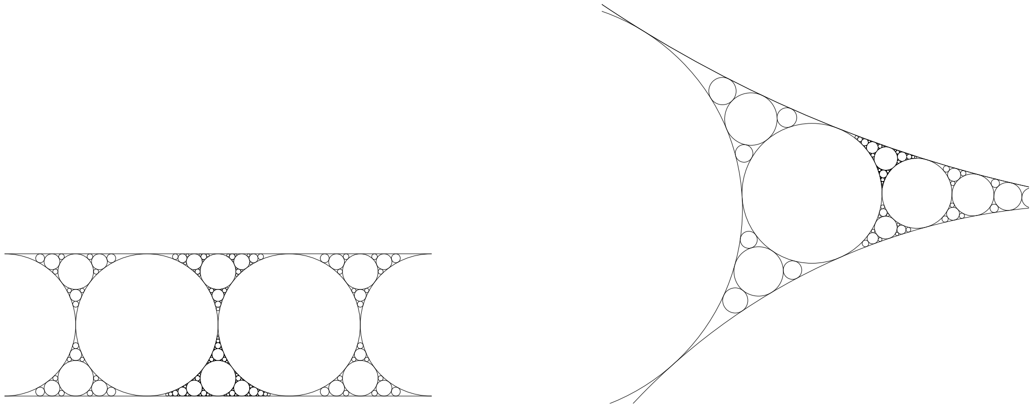

(c) Strip (d) Unbounded

Figure 1. Different boundedness properties for Apollonian packings.

necessarily a composite of a translation, rotation, reflection and/or dilation. Two packings $\mathcal{P}$ and $\mathcal{P}'$ are **similar** if there is a similarity of the plane that takes circles in $\mathcal{P}$ bijectively to circles in $\mathcal{P}'$. There may be non-trivial similarities from $\mathcal{P}$ to itself, in which case we say that $\mathcal{P}$ is **self-similar** and we consider its group of self-similarities.

Our results concern only half-plane packings. We give a necessary and sufficient condition for two such packings to be similar, a classification of the self-similar packings, and a description of all of the self-similarity groups.

We also answer the more specific question of whether two packings are similar via an **orientation-preserving** similarity (that is, one with positive determinant) or via an **orientation-reversing** similarity. Our classification tells us which packings possess an orientation-reversing self-similarity.

We observe that any half-plane packing $\mathcal{P}$ is similar to a packing containing three circles in the configuration shown in Figure 2, where L is the x-axis, and α^2 and 1 refer to the curvatures of the circles they label. For $\alpha > 0$, we define $\mathcal{P}_\alpha$ to be the unique Apollonian circle packing containing that configuration. Because each half-plane packing $\mathcal{P}$ is similar to such a packing, we restrict our attention to studying the packings $\mathcal{P}_\alpha$, and we state results in terms of this particular class of packings.

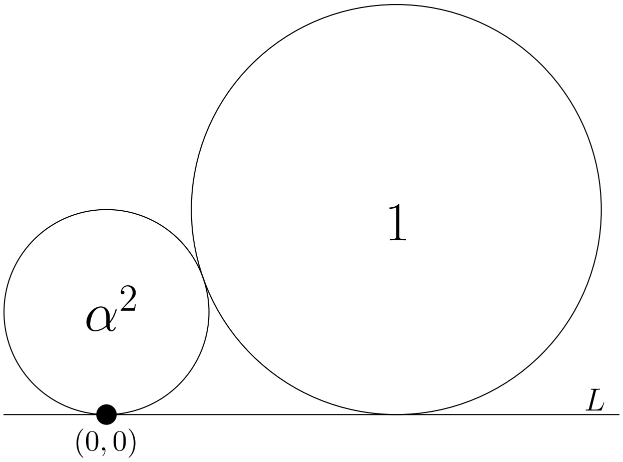

FIGURE 2. The generating triple for the packing $\mathcal{P}_\alpha$. The numbers α^2 and 1 represent the curvatures of their respective circles, and L is a straight line.

Our first result relates similarities of half-plane packings to elements of the projective general linear group $\mathrm{PGL}_2(\mathbb{Z})$. This is the quotient of the group $\mathrm{GL}_2(\mathbb{Z})$ of invertible 2×2-matrices with integer entries, by the subgroup consisting of $\pm I$ where I denotes the identity matrix. We also refer to the subgroup $\mathrm{PSL}_2(\mathbb{Z})$ consisting of those elements whose underlying matrices have determinant 1.

Theorem 1.1. *Let $\alpha, \beta > 0$. There is a bijection between the set of similarities preserving the x-axis that map $\mathcal{P}_\beta$ to $\mathcal{P}_\alpha$ and the set of elements $\begin{bmatrix} a & b \\ c & d \end{bmatrix} \in \mathrm{PGL}_2(\mathbb{Z})$ such that $\frac{a\alpha+b}{c\alpha+d} = \beta$. The similarity is orientation-preserving if and only if the corresponding element is in $\mathrm{PSL}_2(\mathbb{Z})$. In particular, $\mathcal{P}_\alpha$ and $\mathcal{P}_\beta$ are similar (resp. similar via an orientation-preserving similarity) if and only if there exist integers a, b, c, d, with $ad - bc = \pm 1$ (resp. $+1$), such that $\frac{a\alpha+b}{c\alpha+d} = \beta$.*

Taking $\alpha = \beta$ in Theorem 1.1 helps us to calculate the self-similarity groups. We write $\mathrm{Symm}(\mathcal{P})$ for the self-similarity group of the packing $\mathcal{P}$, and we write $\mathrm{Symm}_+(\mathcal{P})$ for the subgroup of $\mathrm{Symm}(\mathcal{P})$ consisting of orientation-preserving self-similarities of $\mathcal{P}$.

Theorem 1.2. *Let $\alpha > 0$. Then:*

(i) *If $\alpha \in \mathbb{Q}$, then $\mathcal{P}_\alpha$ is a strip packing and*

$$\mathrm{Symm}(\mathcal{P}_\alpha) \cong D_\infty \times \mathbb{Z}/2\mathbb{Z}$$

with subgroup

$$\mathrm{Symm}_+(\mathcal{P}_\alpha) \cong D_\infty,$$

where D_∞ denotes the infinite dihedral group.

(ii) *If α is quadratic over $\mathbb{Q}$, then*

$$\mathrm{Symm}(\mathcal{P}_\alpha) \cong \mathbb{Z}.$$

Let D denote the discriminant of the primitive integral polynomial with root α. Then the subgroup $\mathrm{Symm}_+(\mathcal{P}_\alpha)$ is:

- *equal to $\mathrm{Symm}(\mathcal{P}_\alpha)$ if the Pell equation $x^2 - Dy^2 = -4$ has no integral solution for (x, y);*
- *the index 2 subgroup of $\mathrm{Symm}(\mathcal{P}_\alpha)$ if $x^2 - Dy^2 = -4$ does have an integral solution.*

(iii) *Otherwise*

$$\mathrm{Symm}_+(\mathcal{P}_\alpha) = \mathrm{Symm}(\mathcal{P}_\alpha) = 1.$$

As one might imagine from the form of Theorem 1.2, there is a striking connection between half-plane Apollonian circle packings and continued fractions which we describe in §4. In particular, we have the following result.

Theorem 1.3. *Let $\alpha, \beta > 0$. The packings $\mathcal{P}_\alpha$ and $\mathcal{P}_\beta$ are similar if and only if the continued fraction expansions of α and β are eventually equal (that is, become equal when initial segments, of possibly different lengths, are removed from each).*

As a consequence of Theorem 1.3, we obtain the following classification of the self-similar half-plane packings.

Theorem 1.4. *The similarity classes of self-similar half-plane (non-strip) packings are in a one-to-one correspondence with the finite, non-repeating sequences of positive integers, up to cyclic permutations. Here "non-repeating" means that the sequence cannot be realized as a concatenation of multiple copies of a shorter sequence. The packing has an orientation-reversing self-similarity if and only if the length of the corresponding sequence is odd.*

In Figures 11-14, at the end of the paper, we show the self-similar half-plane packings corresponding to the sequences (1), (2), (3) and $(1,2)$.

Here is a quick outline. In §2 we give a precise definition of Apollonian circle packings and establish some of their basic properties. The main part of that section is then to describe a labelling system for the circles in a half-plane packing that are tangent to the x-axis. In §3 we relate those labels to the curvatures of the circles and use this relationship to prove Theorems 1.1 and 1.2. In §4 we examine the connection between half-plane packings and continued fractions which we use to prove Theorems 1.3 and 1.4.

Acknowledgements. This project grew out of a VIGRE Research Group on Circle Packings at the University of Georgia. We would like to thank Sa'ar Hersonsky for his leadership and suggestions. Michael Berglund and Chris Pryby were involved in much of the initial work on the project and this paper would not exist without their contributions. Other members of the group, including Jennifer Ellis, provided useful feedback.

2. Apollonian Circle Packings

We begin with a precise definition of an Apollonian circle packing.

Definition 2.1. For the purposes of this paper, a **circle** in $\mathbb{R}^2$ is either a circle or a straight line, together with a choice of one of the components of its complement which we refer to as the **interior** of the circle. Note that what we refer to as the interior of a circle may be the unbounded component of its complement, and the interior of a straight line is one of the two half-planes it determines.

An **Apollonian circle packing** is a collection $\mathcal{P}$ of circles in $\mathbb{R}^2$ with disjoint interiors such that

(i) there exists a set of three mutually tangent circles in $\mathcal{P}$;
(ii) if a circle C is tangent to three mutually tangent circles that are in $\mathcal{P}$, then C is also in $\mathcal{P}$.

An Apollonian circle packing can be constructed recursively in the following way.

Definition 2.2. Let $\mathcal{P}^{(0)}$ be a set of three mutually tangent circles in $\mathbb{R}^2$ with disjoint interiors. Given $\mathcal{P}^{(n)}$, we define $\mathcal{P}^{(n+1)}$ to be the set of circles in $\mathbb{R}^2$ consisting of $\mathcal{P}^{(n)}$ together with any circle that is tangent to three mutually tangent circles in $\mathcal{P}^{(n)}$. The **Apollonian circle packing generated by** $\mathcal{P}^{(0)}$ is

$$\mathcal{P} := \bigcup_{n=0}^{\infty} \mathcal{P}^{(n)}$$

It can be seen by an induction argument that the circles in $\mathcal{P}^{(n)}$ have disjoint interiors, and it follows that $\mathcal{P}$ is an Apollonian circle packing in the sense of Definition 2.1.

Note that a theorem of Apollonius says that for three mutually tangent circles in $\mathbb{R}^2$ with disjoint interiors, there are precisely two other circles tangent to all three. Each of these two circles lies in an *interstice* formed by the original three circles. See Figure 3.

Definition 2.3. Let A, B, and C be three mutually tangent circles in $\mathbb{R}^2$ with disjoint interiors. The complement of $A \cup B \cup C$ in $\mathbb{R}^2$ consists of five components — three of the components are the interiors of the respective circles, and the other two are called the **interstices** formed by A, B, and C.

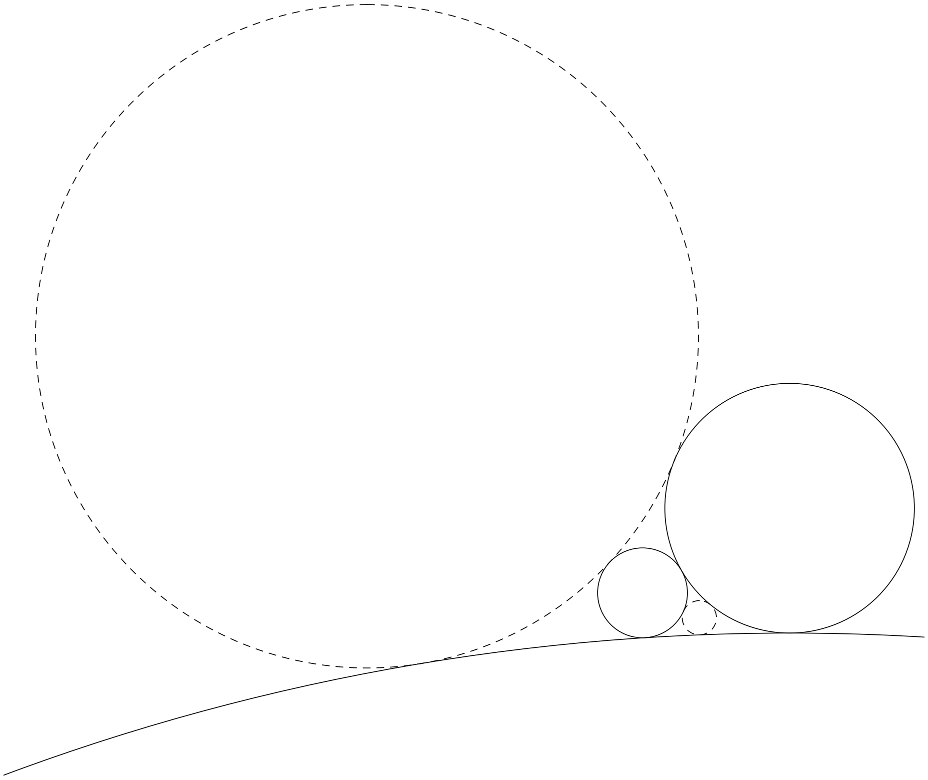

FIGURE 3. The two dashed circles lie in the interstices bounded by the three solid circles.

Lemma 2.4. *Let $\mathcal{P}$ be an Apollonian circle packing. Then $\mathcal{P}$ is generated, in the sense of Definition 2.2, by any set of three mutually tangent circles in $\mathcal{P}$.*

Corollary 2.5. *If $\mathcal{P}$ and $\mathcal{P}'$ are two Apollonian packings with a common triple of mutually tangent circles, then $\mathcal{P} = \mathcal{P}'$.* □

Proof of Lemma 2.4. Since $\mathcal{P}$ certainly contains the packing generated by any set $\mathcal{P}^{(0)}$ of three mutually tangent circles, it is sufficient to show that there is no room for any other circles. In particular, this will be true if the complement of the set of interiors of circles in $\mathcal{P}$ (called the *residual set of* $\mathcal{P}$) has Lebesgue measure zero. A proof of this fact may be found in [4, Theorem 4.2]. □

For us, the point of the recursive construction of Apollonian circle packings is that some of our arguments proceed by induction on the stage at which the circles are created in this process. We therefore make the following definition.

Definition 2.6. Fix a generating triple $\mathcal{P}^{(0)}$ for the packing $\mathcal{P}$. The **generation** of a circle $C \in \mathcal{P}$ (with respect to $\mathcal{P}^{(0)}$), denoted by $\operatorname{gen}(C)$, is the unique $n \in \mathbb{Z}_{\geq 0}$ such that $C \in \mathcal{P}^{(n)} \setminus \mathcal{P}^{(n-1)}$.

We now narrow our focus to half-plane packings. Let $\mathcal{P}$ be a half-plane packing, that is, a packing that contains at least one line L. We assume that L coincides with the x-axis and that the remaining circles in $\mathcal{P}$ are in the upper-half plane. (Any half-plane packing is similar to one that satisfies this condition.)

Definition 2.7. Most of our analysis of half-plane packings can be done be focusing on the circles in $\mathcal{P}$ that are tangent to the line L. We define

$$\mathcal{P}_L := \{C \in \mathcal{P} \mid C \text{ is tangent to } L\}.$$

The 'mutually disjoint interiors' requirement of circle packings ensures that no two circles in $\mathcal{P}_L$ may be tangent to L at the same point. This property allows us to define a total ordering on the set $\mathcal{P}_L$. We say that C *is to the left of* C', or $C \prec C'$, if the x-coordinate of the point of tangency between C and L is *less than* the x-coordinate of the point of tangency between C' and L. In the case that $\mathcal{P}$ is a strip packing, with L' the line in $\mathcal{P}$ which is parallel to L, we consider the x-coordinate of the point of tangency between L' and L to be $-\infty$; in other words, $L' \prec C$ for all $C \in \mathcal{P}_L$ with $C \neq L'$.

Definition 2.8. Let X and Y be two tangent circles in $\mathcal{P}_L$, neither of which is a line. Then $\{X, Y, L\}$ is a triple of mutually tangent circles in $\mathbb{R}^2$ and therefore determines two interstices in the plane. One interstice is bounded, and the other is unbounded; we refer to these as the **bounded interstice for X and Y** and the **unbounded interstice for X and Y** respectively. We say that the circle C **fills the bounded** (resp. **unbounded**) **interstice for X and Y** if C is the unique circle in the bounded (resp. unbounded) interstice for X and Y which is tangent to X, Y, and L (see Figure 4). Note that, by Definition 2.1, C necessarily lies in $\mathcal{P}$ and hence also $\mathcal{P}_L$.

Remark. No circle in the bounded interstice for X and Y can be tangent to a circle in the unbounded interstice for X and Y.

By Lemma 2.4 we can view the packing $\mathcal{P}$ as generated by the triple $\{X, Y, L\}$ in the sense of Definition 2.2 for any pair of tangent circles $X, Y \in \mathcal{P}_L$. For the remainder of this section, we fix a choice of X and Y and assume that $X \prec Y$, that is, X is to the left of Y. We also assume that X and Y are actual circles, i.e., neither is a line.

Definition 2.9. It is convenient to divide up the circles in $\mathcal{P}_L$ according to which interstice they are contained in. We define

$$\mathcal{P}_L^+ = \{C \in \mathcal{P}_L \mid X \preceq C \preceq Y\}$$

and

$$\mathcal{P}_L^- = \{C \in \mathcal{P}_L \mid C \preceq X \text{ or } Y \preceq C\}.$$

Geometrically, $\mathcal{P}_L^+$ consists of X, Y, and those circles in $\mathcal{P}_L$ that are in the bounded interstice for X and Y, while $\mathcal{P}_L^-$ consists of X, Y, and those circles in $\mathcal{P}_L$ that are in the unbounded interstice for X and Y. Note that $\mathcal{P}_L^+ \cup \mathcal{P}_L^- = \mathcal{P}_L$ and $\mathcal{P}_L^+ \cap \mathcal{P}_L^- = \{X, Y\}$.

Lemma 2.10. *For each circle $C \in \mathcal{P}_L^+$, $C \notin \{X, Y\}$, there exist circles A and B in $\mathcal{P}_L^+$ of generation strictly less than that of C such that C fills the bounded interstice for A and B. (Recall that the generation of a circle in a packing $\mathcal{P}$ depends on a choice of generating triple; in this case, $\mathcal{P}^{(0)} = \{X, Y, L\}$.)*

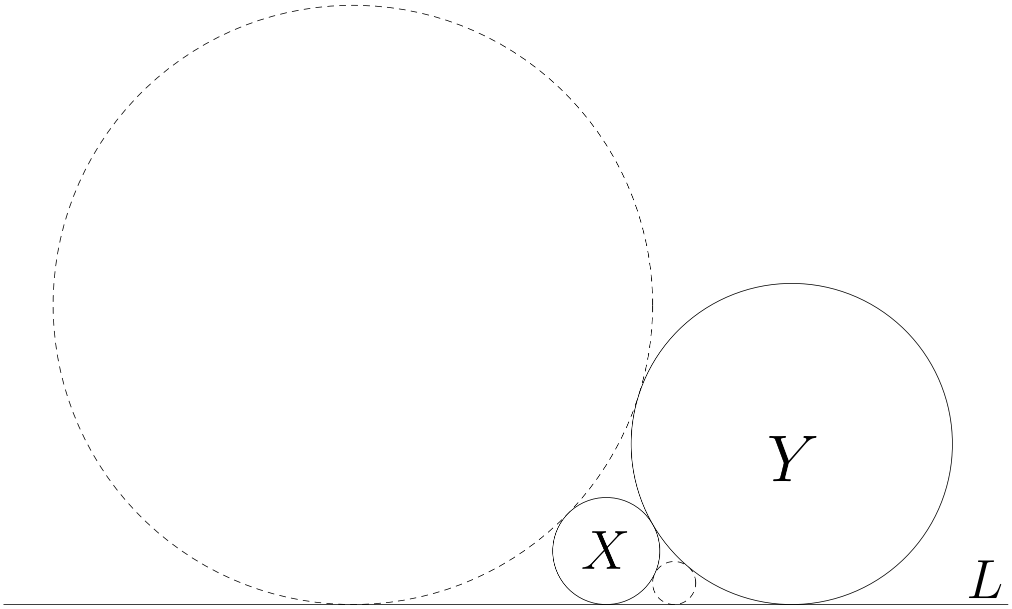

Figure 4. The larger dashed circle fills the unbounded interstice for X and Y, and the smaller dashed circle fills the bounded interstice.

Proof. We work by induction on the generation of C. If C is generation 1, then it must be *the* circle that fills the bounded interstice between X and Y, so satisfies the lemma. Now suppose that $\mathrm{gen}(C) \geq 2$. Thinking about when the circle C is added to the packing in the recursive construction of Definition 2.2, we see that there are exactly three mutually tangent circles of generation less than C that are tangent to C. One of these circles must be the line L, so let A and B be the other two. It follows from the remark after Definition 2.8 that $A, B \in \mathcal{P}_L^+$. Now C fills one of the interstices formed by A and B. We need to show that it fills the bounded interstice.

Now exactly one of A and B must be of generation exactly one less than C. (To see this, we recall the procedure for recursively building an Apollonian packing. This procedure implies that as soon as A and B have been added, the circle C will be added in the very next generation. On the other hand, no two circles of the same generation are tangent since they fill different interstices.) Suppose this is B, so that we have $\mathrm{gen}(C) > \mathrm{gen}(B) \geq 1$. By the induction hypothesis, B fills the bounded interstice formed by two other circles of generation less than it. One of those must be A and let the other be D. But now we see that D fills the unbounded interstice for A and B. Since $\mathrm{gen}(D) < \mathrm{gen}(B) < \mathrm{gen}(C)$, we cannot have $D = C$. It follows then that C must fill the bounded interstice for A and B. ☐

Our main tool for keeping track of the circles in a half-plane packing $\mathcal{P}$ is a labelling for each circle in $\mathcal{P}_L$ by a pair of integers (a, b). The remainder of this section is devoted to the construction and properties of this labelling. In §3 we relate this labelling to the curvatures of the circles in $\mathcal{P}$ and use it to deduce information about similarities between different packings.

Definition 2.11. We define a labelling function $\vec{x} = (x, y) : \mathcal{P}_L \to \mathbb{Z}^2$. We define the labelling recursively starting with $\vec{x}(X) = (1, 0)$ and $\vec{x}(Y) = (0, 1)$. For the remaining circles in $\mathcal{P}_L$, the label is determined by the following rule:

If C is the circle that fills the bounded interstice for A and B, then

$$\vec{x}(C) = \vec{x}(A) + \vec{x}(B). \tag{2.12}$$

At each stage of the construction of the packing from its generators, X, Y, and L, a new circle in $\mathcal{P}_L$ fills either the bounded or unbounded interstice for a pair of circles already present. The equation above determines a label for each such new circle — see Figure 5.

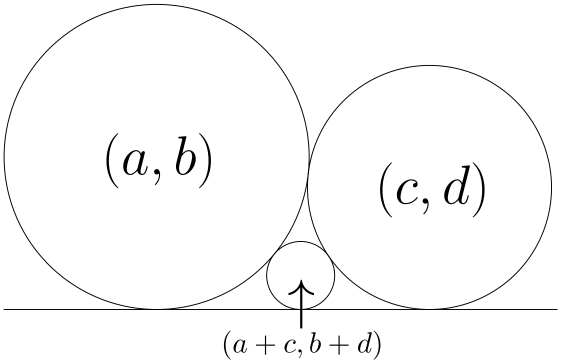

FIGURE 5. The relationship satisfied by the labels.

The main result of this section, Proposition 2.21, tells us that the labelling function $\vec{x}$ is one-to-one and that for each pair (a, b) of coprime integers, exactly one of (a, b) and $(-a, -b)$ is in the image of $\vec{x}$. It also gives us a necessary and sufficient condition on the labels for two circles in $\mathcal{P}_L$ to be tangent. It is convenient to start with this condition, which is stated in terms of the matrix formed by the labels of the two circles.

Lemma 2.13. *Let A and B be a pair of tangent circles in $\mathcal{P}_L$ such that $A \prec B$. Then*

$$\begin{vmatrix} x(A) & y(A) \\ x(B) & y(B) \end{vmatrix} = 1.$$

Proof. The proof will be by induction on $\operatorname{gen}\{A, B\} := \max\{\operatorname{gen}(A), \operatorname{gen}(B)\}$. The base case is immediate: the generation zero circles form the pair $\{X, Y\}$, which are labeled $(1, 0)$ and $(0, 1)$ respectively. The corresponding matrix is the identity, which has determinant 1.

Now suppose $\operatorname{gen}\{A, B\} = n \geq 1$. First, observe that we cannot have $\operatorname{gen}(A) = \operatorname{gen}(B) = n$: if $\operatorname{gen}(A) = \operatorname{gen}(B) = n \geq 1$, then A and B were constructed to fill two *disjoint* interstices in $\mathcal{P}^{(n-1)}$ and cannot therefore be tangent. Hence $\{A, B\}$ contains a unique circle of generation n. Furthermore, because a generation n circle is constructed to fill a single interstice in $\mathcal{P}^{(n-1)}$, it is necessarily tangent to exactly three circles of generation strictly less than n. Therefore the circle of generation n (either A or B) is tangent to L (generation zero), the circle in $\{A, B\}$ of smaller generation, and a third circle C of generation strictly less than n.

There are three possibilities for the position of C relative to A and B: C can be to the left of both, to the right of both, or between the two. Moreover, the generation n circle can be either A or B, so there are a total of six cases to consider. We only give the proof in two cases — the other four are nearly identical. To prove them, we use the fact that the matrix row operations of row addition/subtraction are determinant-preserving and that switching two rows switches the sign of the determinant. In each case, the final equality holds by the induction hypothesis.

Case 1: Suppose $C \prec A \prec B$ and $\operatorname{gen}(B) = n$. Then $\vec{x}(A) = \vec{x}(C) + \vec{x}(B)$, and

$$\begin{vmatrix} \vec{x}(A) \\ \vec{x}(B) \end{vmatrix} = \begin{vmatrix} \vec{x}(A) \\ \vec{x}(B) - \vec{x}(A) \end{vmatrix} = \begin{vmatrix} \vec{x}(A) \\ -\vec{x}(C) \end{vmatrix} = - \begin{vmatrix} -\vec{x}(C) \\ \vec{x}(A) \end{vmatrix} = \begin{vmatrix} \vec{x}(C) \\ \vec{x}(A) \end{vmatrix} = 1.$$

Case 2: Suppose $A \prec C \prec B$ and $\mathrm{gen}(A) = n$. Then $\vec{x}(C) = \vec{x}(A) + \vec{x}(B)$, and

$$\begin{vmatrix} \vec{x}(A) \\ \vec{x}(B) \end{vmatrix} = \begin{vmatrix} \vec{x}(A) + \vec{x}(B) \\ \vec{x}(B) \end{vmatrix} = \begin{vmatrix} \vec{x}(C) \\ \vec{x}(B) \end{vmatrix} = 1.$$

☐

Corollary 2.14. *For any circle $C \in \mathcal{P}_L$, $\gcd(x(C), y(C)) = 1$.* ☐

We now begin the proof that our labelling function $\vec{x}$ is one-to-one. We do this first for those circles in the bounded interstice for X and Y.

Lemma 2.15. *For $C \in \mathcal{P}_L^+$ we have $x(C), y(C) \geq 0$ with equality only if either $C = X$ or $C = Y$.*

Proof. This follows from (2.12) by induction on generation since, by Lemma 2.10, the circle C fills the bounded interstice of two circles of strictly smaller generation than it. ☐

Corollary 2.14 and Lemma 2.15 tell us that every circle in $\mathcal{P}_L^+$ is labeled by a pair of nonnegative coprime integers. We now prove that every such pair is the label of a unique circle in $\mathcal{P}_L^+$. At the same time, we prove the converse of Lemma 2.13 for $\mathcal{P}_L^+$ — that if circles $A, B \in \mathcal{P}_L^+$ have the determinant of the matrix formed by their labels equal to 1, then they are tangent with $A \prec B$. We first need the following elementary lemma.

Lemma 2.16. *Let a and b be positive, coprime integers. Then there exist unique integers u and v that satisfy the following properties:*

(i) $au - bv = 1$,
(ii) $0 < u \leq b$, *and*
(iii) $0 \leq v < a$.

Proof. Because a, b are coprime, we can find an integer solution (x, y) to the equation

$$ax - by = 1. \tag{2.17}$$

Given a particular solution (x_0, y_0) to (2.17), the entire solution set is

$$\{(x, y) = (x_0 + kb, y_0 + ka) \ : \ k \in \mathbb{Z}\}.$$

There is then a *unique* $k \in \mathbb{Z}$ such that $0 < x_0 + kb \leq b$. Let $u := x_0 + kb$. Then u satisfies property (ii). Setting $v := y_0 + ka$, property (i) is also satisfied, and property (iii) is a consequence of properties (i) and (ii). ☐

Lemma 2.18. *Let $\begin{bmatrix} a & b \\ c & d \end{bmatrix}$ be a determinant 1 matrix with nonnegative integer coefficients. Then there exist unique circles $C, C' \in \mathcal{P}_L^+$ such that $\vec{x}(C) = (a, b)$ and $\vec{x}(C') = (c, d)$. Moreover, C and C' are tangent with $C \prec C'$.*

Proof. The proof is by induction on $\max\{a + b, c + d\}$. If $\max\{a + b, c + d\} = 1$, then necessarily $\begin{bmatrix} a & b \\ c & d \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$. By Lemma 2.15, all circles in $\mathcal{P}_L^+$ different from X and Y must have $a + b > 1$, so there can be no circles in $\mathcal{P}_L^+$, other than X and Y, labeled by the pairs $(1, 0)$ and $(0, 1)$.

Once we have proved that there is a unique circle with label (a, b), we denote that circle by $C_{(a,b)}$. At this point, therefore, we can write $C_{(1,0)} = X$ and $C_{(0,1)} = Y$.

Now fix an integer $n > 1$ and suppose we have proved the lemma, and hence constructed the circles $C_{(a,b)}$ and $C_{(c,d)}$, for any a, b, c, d as in the statement of the lemma with $\max\{a + b, c + d\} < n$. We then take a, b, c, d with $\max\{a + b, c + d\} = n$.

First of all, if $a+b=c+d$, then

$$(d-b)(a+b)=d(a+b)-b(a+b)=d(a+b)-b(c+d)=ad-bc=1,$$

so $a+b=1$. Since $c+d=a+b$, this contradicts the fact that $\max\{a+b,c+d\}>1$. Therefore $a+b>c+d$ or $a+b<c+d$. We prove the lemma in the case where $a+b>c+d$; the proof of the other case is virtually identical.

Since $a+b>c+d\geq 1$, we have that $0<(a-c)+(b-d)<a+b$ and therefore

$$\max\{(a-c)+(b-d),c+d\}<a+b=n.$$

Since $\begin{vmatrix} a-c & b-d \\ c & d \end{vmatrix}=1$, we may apply the induction hypothesis to the matrix $\begin{bmatrix} a-c & b-d \\ c & d \end{bmatrix}$ once we show that $a-c\geq 0$ and $b-d\geq 0$.

Since $a+b>c+d$, we must have $a>c$ or $b>d$. If $a>c>0$, then

$$\begin{aligned} ad-bc=1 &\implies cd<bc+1 \\ &\implies cd\leq bc \\ &\implies d\leq b. \end{aligned}$$

(Note that if $c=0$, then $ad=1$, so $a=d=1$. It follows from the fact that $a+b>c+d$ that $b\geq 1=d$.) A similar argument shows that if we assume instead that $b>d$, then also $a>c$.

Therefore we have both $a-c\geq 0$ and $b-d\geq 0$, so the induction hypothesis tells us that there are unique circles $C_{(a-c,b-d)}$ and $C_{(c,d)}$ in $\mathcal{P}_L^+$ satisfying $\vec{x}\left(C_{(a-c,b-d)}\right)=(a-c,b-d)$ and $\vec{x}\left(C_{(c,d)}\right)=(c,d)$ and that, moreover, these circles are tangent with $C_{(a-c,b-d)}\prec C_{(c,d)}$.

Now let C be the circle that fills the bounded interstice for $C_{(a-c,b-d)}$ and $C_{(c,d)}$. Then

$$\vec{x}(C)=(a-c,b-d)+(c,d)=(a,b),$$

so there exists a circle C labeled by the pair (a,b), which by construction is to the left of and tangent to $C_{(c,d)}$.

Finally, we must show that C is the *only* circle in $\mathcal{P}_L^+$ that satisfies $\vec{x}(C)=(a,b)$. Suppose C'' is a circle with $\vec{x}(C'')=(a,b)$. By Lemma 2.10, C'' fills the bounded interstice for two circles $A\prec B$ in $\mathcal{P}_L^+$. Then $\vec{x}(C'')=\vec{x}(A)+\vec{x}(B)$, so we can write $\vec{x}(B)=(v,u)$ and $\vec{x}(A)=(a-v,b-u)$. Since A and B are both in $\mathcal{P}_L^+$, Lemma 2.15 tells us that each of v, u, $a-v$, and $b-u$ is nonnegative, so $0\leq u\leq b$ and $0\leq v\leq a$.

In fact, we have $u\neq 0$ and $v\neq a$. Indeed, we know from Lemma 2.15 that if $u=0$, then $B=X$. However, $X\preceq C''\prec B$, so $X\neq B$, which means we cannot have $u=0$. A similar argument shows that $a-v\neq 0$, and so $v\neq a$. Therefore $0<u\leq b$ and $0\leq v<a$, which are precisely properties (ii) and (iii) from Lemma 2.16. That property (i) is satisfied follows by Lemma 2.13, since C'' is tangent to and to the left of B. Since these three properties uniquely determine (v,u), and since (c,d) satisfies these three conditions by construction, we conclude that $v=c$ and $u=d$. By the uniqueness of $C_{(a-c,b-d)}$ and $C_{(c,d)}$ following from the induction hypothesis, we have $A=C_{(a-c,b-d)}$ and $B=C_{(c,d)}$, and it follows therefore that $C''=C$. □

Definition 2.19. Lemma 2.18 implies that any pair (a,b) of coprime nonnegative integers is the label of a unique circle in $\mathcal{P}_L^+$. As in the proof of Lemma 2.18, we denote that circle by $C_{(a,b)}$.

Lemma 2.18 yields a complete understanding of the labels of circles in $\mathcal{P}_L^+$. We now define an operation which maps $\mathcal{P}_L^+$ bijectively onto $\mathcal{P}_L^-$, and use this to relate the labels of circles in $\mathcal{P}_L^-$ to those of circles in $\mathcal{P}_L^+$.

Let $\mathcal{I}$ be the unique circle which contains the three points of tangency among $C_{(1,0)}$, $C_{(0,1)}$, and L. (See Figure 6.) Define the map

$$\iota : \mathbb{R}^2 \cup \{\infty\} \to \mathbb{R}^2 \cup \{\infty\}$$

to be inversion with respect to $\mathcal{I}$.

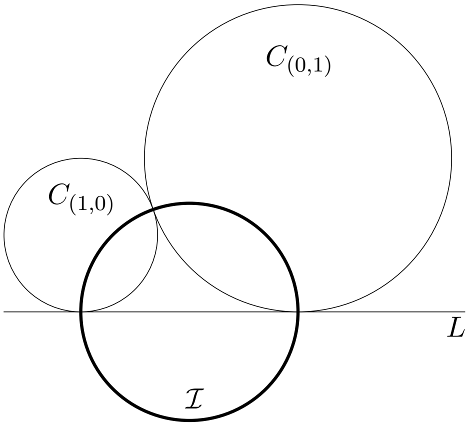

FIGURE 6. The inversion circle $\mathcal{I}$

Let us pause to mention some of the relevant properties of $\mathcal{I}$ and the map ι. Note that when we say that ι *fixes* a particular circle or set of circles, we mean only as sets in $\mathbb{R}^2$, not pointwise.

(i) Inversion with respect to a circle is a bijection of order two; i.e., $\iota \circ \iota = \text{id}$.
(ii) Since inversion maps circles to circles, ι maps Apollonian packings to Apollonian packings.
(iii) The inversion circle $\mathcal{I}$ intersects each of $C_{(1,0)}$, $C_{(0,1)}$, and L orthogonally, and therefore ι fixes each of these three circles. Therefore, by property (ii) and Corollary 2.5, ι fixes $\mathcal{P}$ (and hence $\mathcal{P}_L$ since L is fixed).
(iv) The interior of $\mathcal{I}$ contains the bounded interstice for $C_{(1,0)}$ and $C_{(0,1)}$, and the exterior of $\mathcal{I}$ contains the unbounded interstice for $C_{(1,0)}$ and $C_{(0,1)}$. Since ι maps the interior of $\mathcal{I}$ to the exterior of $\mathcal{I}$, and vice versa, and since property (iii) holds, it follows that ι maps $\mathcal{P}_L^+$ to $\mathcal{P}_L^-$ and vice versa.

Because $\mathcal{I}$ intersects L orthogonally, the center of $\mathcal{I}$ lies on L. It makes sense, then, to talk about a circle $C \in \mathcal{P}_L$ lying to the left or right of $\mathcal{I}$, by which we mean that the point of tangency of C with L lies to the left or right of the *center* of $\mathcal{I}$. We now record two more properties of $\mathcal{I}$ and ι:

(v) If C lies to the left (resp. right) of $\mathcal{I}$, then $\iota(C)$ also lies to the left (resp. right) of $\mathcal{I}$. Furthermore, if $C \prec C'$ both lie to the left (resp. right) of $\mathcal{I}$, then $\iota(C') \prec \iota(C)$ both lie to the left (resp. right) of $\mathcal{I}$.
(vi) A circle C contains the center of $\mathcal{I}$ (that is, the point of tangency between C and L is precisely the center of $\mathcal{I}$) if and only if $\iota(C)$ is a line parallel to L.

As mentioned above, the reason for introducing the inversion map ι is to set up a one-to-one correspondence between $\mathcal{P}_L^+$ and $\mathcal{P}_L^-$. The following lemma establishes the connection between the labels of circles in $\mathcal{P}_L^-$ and their images under ι, which lie in $\mathcal{P}_L^+$.

Lemma 2.20. *Let $C \in \mathcal{P}_L^-$, and let $C_{(a,b)} = \iota(C) \in \mathcal{P}_L^+$ be the image of C under the map ι. Then*

$$\vec{x}(C) = \begin{cases} (a, -b) & \text{if } C \preceq C_{(1,0)} \\ (-a, b) & \text{if } C_{(0,1)} \preceq C \end{cases}.$$

Proof. We will prove the statement by induction on the generation of $\iota(C)$. In the case that $\mathrm{gen}(\iota(C)) = 0$, we have $\iota(C) = C_{(1,0)}$ or $\iota(C) = C_{(0,1)}$. Since $C_{(1,0)}$ and $C_{(0,1)}$ are fixed by ι (and since ι is one-to-one), it follows that $C = C_{(1,0)}$ or $C = C_{(0,1)}$, and in both cases the statement holds.

Now suppose $\mathrm{gen}(\iota(C)) \geq 1$. Then, since $C_{(a,b)} = \iota(C)$ lies in $\mathcal{P}_L^+$ and is not equal to $C_{(1,0)}$ or $C_{(0,1)}$, Lemma 2.10 tells us that $C_{(a,b)}$ fills the bounded interstice for two circles $C_{(a_1,b_1)}, C_{(a_2,b_2)} \in \mathcal{P}_L^+$ with $\mathrm{gen}(C_{(a_1,b_1)}), \mathrm{gen}(C_{(a_2,b_2)}) < \mathrm{gen}(C_{(a,b)})$. Assume that $C_{(a_1,b_1)} \prec C_{(a_2,b_2)}$. By definition, $(a_1, b_1) + (a_2, b_2) = (a, b)$. Because inversion preserves tangencies, the circles $A = \iota(C_{(a_1,b_1)})$ and $B = \iota(C_{(a_2,b_2)})$ are tangent to each other as well as to C and L.

There are a total of seven cases to consider, each corresponding to the position of the center of $\mathcal{I}$ with respect to the points of tangency of the circles $C_{(a,b)}$, $C_{(a_1,b_1)}$, and $C_{(a_2,b_2)}$ with the line L — the center of $\mathcal{I}$ could lie on one of the three circles, it could lie between two of the circles, or it could lie to the left or right of all three of the circles. We will prove the result for two example cases; the proofs in the other cases are quite similar.

Case 1: Suppose that the center of $\mathcal{I}$ lies between the points of tangency of $C_{(a_1,b_1)}$ and $C_{(a,b)}$ with L, as shown in Figure 7. In particular, $C_{(a_1,b_1)}$ is to the left of $\mathcal{I}$ and $C_{(a,b)} \prec C_{(a_2,b_2)}$ are to the right. By property (v) above, we may conclude that A is to the left of $\mathcal{I}$ and $B \prec C$ are to the right of $\mathcal{I}$. Because all three of the image circles necessarily lie in $\mathcal{P}_L^-$, it follows that $A \preceq C_{(1,0)}$ and $C_{(0,1)} \preceq B \prec C$.

By the induction hypothesis, we know that $\vec{x}(A) = (a_1, -b_1)$ and $\vec{x}(B) = (-a_2, b_2)$. Since $A \prec B \prec C$, it follows that $\vec{x}(A) + \vec{x}(C) = \vec{x}(B)$, and so

$$\vec{x}(C) = \big(-(a_1 + a_2), b_1 + b_2\big) = (-a, b).$$

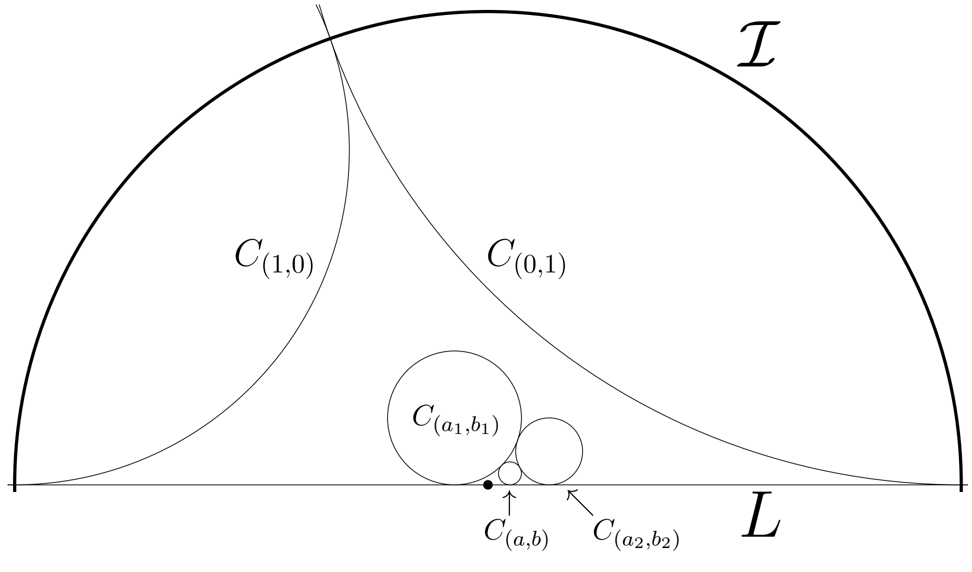

FIGURE 7. Case 1 of Lemma 2.20. The point on the line L represents the center of $\mathcal{I}$.

Case 2: Suppose that the center of $\mathcal{I}$ coincides with the point of tangency between $C_{(a,b)}$ and L, as shown in Figure 8. Then $A \preceq C_{(1,0)}$, $C_{(0,1)} \preceq B$, and C is a line parallel to L, which means that C is to the left of every circle in $\mathcal{P}_L$. By induction, we have $\vec{x}(A) = (a_1, -b_1)$ and $\vec{x}(B) = (-a_2, b_2)$. Since $C \prec A \prec B$, we have $\vec{x}(A) = \vec{x}(C) + \vec{x}(B)$. It follows that

$$\vec{x}(C) = \big(a_1 + a_2, -(b_1 + b_2)\big) = (a, -b).$$

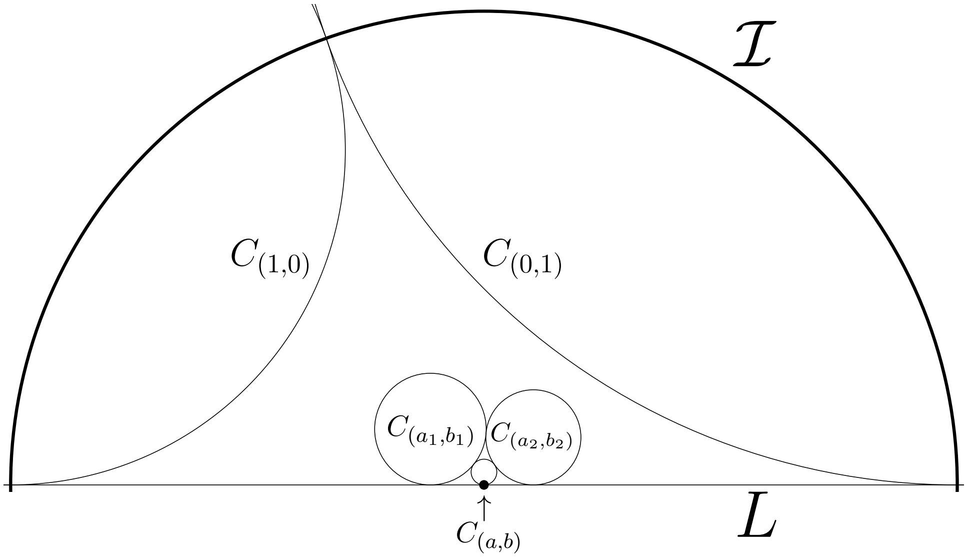

Figure 8. Case 2 of Lemma 2.20

As mentioned above, the proofs of the other five cases are very similar to these two. □

Combining Lemmas 2.18 and 2.20 we obtain a complete understanding of how the circles in $\mathcal{P}_L$ are labeled and when two labeled circles are tangent to one another.

Proposition 2.21. *For any integers a and b with $\gcd(a,b) = 1$, there is either a unique circle in $\mathcal{P}_L$ labelled by (a,b) or a unique circle labelled by $(-a,-b)$, but not both. If a and b are both nonnegative then the label is (a,b). If $C_{(a,b)}$ and $C_{(c,d)}$ are the unique circles in $\mathcal{P}_L$ with labels (a,b) and (c,d) respectively, then $C_{(a,b)}$ is tangent to $C_{(c,d)}$ on the left if and only if $\begin{vmatrix} a & b \\ c & d \end{vmatrix} = 1$.*

Proof. Lemmas 2.18 and 2.20, together with the fact that the inversion operation ι acts as a bijection between $\mathcal{P}_L^+$ and $\mathcal{P}_L^-$, imply the first claim. Lemma 2.13 already tells us the 'only if' part of the second statement. So consider circles $C_{(a,b)}$ and $C_{(c,d)}$ with $ad - bc = 1$. If all a, b, c, d are nonnegative then Lemma 2.18 tells us $C_{(a,b)}$ and $C_{(c,d)}$ are tangent. If one of a, b is negative, then the condition $ad - bc = 1$ implies that one of c, d must be either negative or zero. But then Lemmas 2.18 and 2.20 imply that the circles $\iota(C_{(a,b)})$ and $\iota(C_{(c,d)})$ are tangent. Since ι preserves tangencies, it follows that $C_{(a,b)}$ and $C_{(c,d)}$ are also tangent. That $C_{(a,b)}$ is to the left of $C_{(c,d)}$ follows from Lemma 2.13. □

3. Self-similar half-plane packings

We are now at a point where we may begin to describe the similarities between two half-plane packings. First we recall exactly what is meant by a similarity of $\mathbb{R}^2$.

Definition 3.1. The map $\Phi : \mathbb{R}^2 \to \mathbb{R}^2$ is called a **similarity of** $\mathbb{R}^2$ if there exists some constant $\mu > 0$ such that

$$||\Phi(x) - \Phi(y)|| = \mu||x - y||$$

for all $x, y \in \mathbb{R}^2$.

Every similarity of the plane takes the form

$$\Phi(x) = \mu A x + b,$$

where $\mu > 0$, A is an orthogonal matrix, and $b \in \mathbb{R}^2$. We say that Φ is **orientation-preserving** if $\det A = +1$, and **orientation-reversing** if $\det A = -1$.

The set of similarities of $\mathbb{R}^2$ forms a group under composition, called the **similarity group** of $\mathbb{R}^2$, which we will denote by $\mathcal{S}$. The orientation-preserving similarities form a subgroup $\mathcal{S}_+$.

Similarities take circles to circles, and preserve tangency, so they take Apollonian circle packings to Apollonian circle packings. A key fact about the action of similarities on circle packings is the following.

Lemma 3.2. *Let Φ be a similarity of $\mathbb{R}^2$, and let A, B, and C be three mutually tangent circles with disjoint interiors. If A, B, and C have collinear centers, then Φ is determined by the three circles $\Phi(A)$, $\Phi(B)$, and $\Phi(C)$, up to a reflection in the line on which the centers of the image circles lie. If A, B, and C have non-collinear centers, then Φ is completely determined by $\Phi(A)$, $\Phi(B)$, and $\Phi(C)$.*

Remark. Here we mean that the 'center' of a line L lies 'at infinity' orthogonal to L in the direction of its chosen interior. If one of the circles A, B, C, say A, is a line, then the collinearity condition is satisfied if and only if another of the circles, say B, is also a line, parallel to A. In this case Φ is determined up to a reflection in the line through the center of C that is orthogonal to A and B.

Proof. Since similarities form a group, it suffices to consider the similarities that *fix* A, B, and C (as sets, not pointwise). If a similarity fixes the circles A, B and C, then it fixes their centers. A similarity that fixes three non-collinear points must be the identity. One that fixes three distinct collinear points is either the identity or a reflection in the line formed by them. □

Definition 3.3. For $\Phi \in \mathcal{S}$ and an Apollonian packing $\mathcal{P}$, we write

$$\Phi \cdot \mathcal{P} := \{\Phi(C) : C \in \mathcal{P}\}.$$

Two packings $\mathcal{P}$ and $\mathcal{P}'$ are **similar** if $\mathcal{P}' = \Phi \cdot \mathcal{P}$ for some $\Phi \in \mathcal{S}$.

The group

$$\operatorname{Symm}(\mathcal{P}) := \{\Phi \in \mathcal{S} \mid \Phi \cdot \mathcal{P} = \mathcal{P}\}$$

is the **self-similarity group of** $\mathcal{P}$. The subgroup of $\operatorname{Symm}(\mathcal{P})$ consisting only of orientation-preserving similarities is the **orientation-preserving self-similarity group of** $\mathcal{P}$, denoted by $\operatorname{Symm}_+(\mathcal{P})$. A packing $\mathcal{P}$ is **self-similar** if $\operatorname{Symm}(\mathcal{P})$ is nontrivial.

In order to establish similarity between two packings, we look at the curvatures of the circles involved.

Definition 3.4. For a circle C in $\mathbb{R}^2$, the **curvature** of C, denoted $\operatorname{curv}(C)$ is the reciprocal of the radius of C. A straight line in $\mathbb{R}^2$ is considered to have curvature zero.

Lemma 3.2 allows us to check similarity by looking only at the curvatures in a triple of mutually tangent circles in each packing.

Lemma 3.5. *The packings $\mathcal{P}$ and $\mathcal{P}'$ are similar if and only if they contain triples of mutually tangent circles (A, B, C) and (A', B', C') respectively, such that there exists $\mu > 0$ with*

$$\operatorname{curv}(A') = \mu \operatorname{curv}(A), \ \operatorname{curv}(B') = \mu \operatorname{curv}(B), \ \operatorname{curv}(C') = \mu \operatorname{curv}(C).$$

Proof. If $\mathcal{P}$ and $\mathcal{P}'$ are similar via similarity Φ with scale factor μ, then take any triple (A, B, C) and set $A' = \Phi(A), B' = \Phi(B), C' = \Phi(C)$. To prove the converse, choose a similarity Φ of the plane that takes A to A', B to B' and C to C'. (One can choose a translation composed with dilation to get A to A', add a rotation to get B to B', then add a reflection if necessary to get C to C'.) By construction, $\mathcal{P}'$ and $\Phi(\mathcal{P})$ both contain the triple $\{A', B', C'\}$, and therefore $\mathcal{P}' = \Phi(\mathcal{P})$ by Corollary 2.5. □

Turning now to half-plane packings, that is, those that have a straight line for at least one of the circles, recall that we can focus on the following packings.

Definition 3.6. Let $\alpha \in \mathbb{R}$. Then $\mathcal{P}_\alpha$ is the packing generated by a triple $\{X, Y, L\}$ of mutually tangent circles, where L is the x-axis in $\mathbb{R}^2$, X is a circle of curvature α^2 tangent to (and above) L at the origin, and Y is a circle of curvature 1 resting on L and tangent (on the right) to X. The generating triple is illustrated in Figure 2 in the Introduction.

Lemma 3.7. *Every half-plane packing $\mathcal{P}$ is similar, via an orientation-preserving similarity, to $\mathcal{P}_\alpha$ for some $\alpha > 0$.*

Proof. Choose any two tangent circles in $\mathcal{P}$ that are tangent to a line but are not themselves lines. Taking μ to be the ratio of their curvatures (in the appropriate order), and $\alpha = \sqrt{\mu}$, this follows from 3.5. ☐

The key to analyzing Apollonian circle packings is the following result, due to Descartes. This describes the relationship between the curvatures of four mutually tangent circles in the plane. A selection of proofs of this are given in [9].

Theorem 3.8 (Descartes' Circle Theorem)**.** *Let w, x, y, and z represent the curvatures of four mutually tangent circles in the Euclidean plane. Then*

$$2(w^2 + x^2 + y^2 + z^2) = (w + x + y + z)^2.$$

For half-plane packings, we apply this Theorem in the case where one of the four circles is a line, i.e. has zero curvature. In this case, the quadratic relationship boils down to a linear relationship between the square roots of the curvatures of the circles.

Corollary 3.9. *Let α^2, β^2 and γ^2 represent the curvatures of three mutually tangent circles all tangent to a line L, where $\alpha \geq \beta \geq 0$ and $\gamma \geq 0$. Then*

$$\gamma = \alpha \pm \beta. \tag{3.10}$$

In particular, when the circle of curvature γ^2 lies in the bounded interstice formed by the others, we have

$$\gamma = \alpha + \beta. \tag{3.11}$$

Proof. The proof of (3.10) follows from Theorem 3.8 by setting $w = 0$ and applying the quadratic formula appropriately. The proof of (3.11) follows from the fact that the circle in the bounded interstice has a curvature at least as large as that of the two circles surrounding it. ☐

We illustrate (3.11) in Figure 9. Observe that the illustration of this equation is virtually identical to the illustration in Figure 5, which shows the recursive labelling process defined in §2. It is this linear relationship between the curvatures of tangent circles that inspires that labelling. A key consequence of this connection is Lemma 3.13 below.

Definition 3.12. Consider now some fixed real number $\alpha > 0$ and recall the packing $\mathcal{P}_\alpha$ from Definition 3.6. We write $\mathcal{P}_{\alpha,L}$ for the set of circles in $\mathcal{P}_\alpha$ that are tangent to the line L (that is, the x-axis). Let $C_{(1,0)}$ denote the circle of curvature α^2 that is tangent to L at the origin, and let $C_{(0,1)}$ denote the circle of curvature 1 that is tangent to L and to $C_{(1,0)}$ on the right. As described in the previous section, these choices determine a unique label for each circle in $\mathcal{P}_{\alpha,L}$. When we need to specify the underlying α we use a superscript, as in $C^\alpha_{(a,b)}$, but we often drop the α when context allows.

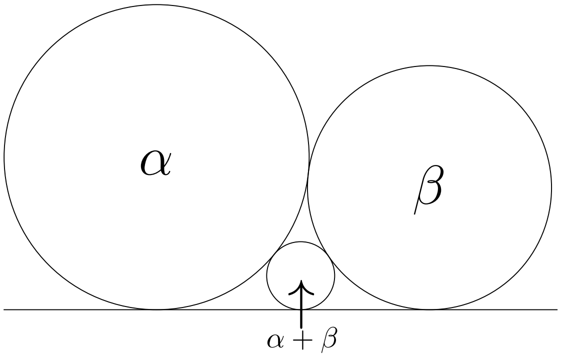

FIGURE 9. Descartes' Theorem for circles in $\mathcal{P}_L$. The circles are labelled by the *square roots* of their curvatures.

Lemma 3.13. *The circle $C_{(a,b)}$ in $\mathcal{P}_{\alpha,L}$ has curvature given by*

$$\operatorname{curv}(C_{(a,b)}) = (a\alpha + b)^2.$$

Moreover, $a\alpha + b \geq 0$.

Proof. The proof is by induction on the generation of $C_{(a,b)}$ with respect to the generating triple $\{L, C_{(1,0)}, C_{(0,1)}\}$. The result is immediately seen to hold for the generation zero circles $C_{(1,0)}$ and $C_{(0,1)}$, since they were chosen to satisfy $\operatorname{curv}(C_{(1,0)}) = \alpha^2$ and $\operatorname{curv}(C_{(0,1)}) = 1$.

Now suppose $\operatorname{gen}(C_{(a,b)}) = n \geq 1$. The circle $C_{(a,b)}$ was constructed to fill an interstice bounded by three circles of generation strictly less than n; since $C_{(a,b)}$ is tangent to L, L is necessarily one of those circles. Because the other two circles are tangent to L as well, we can call them $C_{(a_1,b_1)}$ and $C_{(a_2,b_2)}$ with $C_{(a_1,b_1)} \prec C_{(a_2,b_2)}$. We do the case where $C_{(a,b)}$ fills the bounded interstice between $C_{(a_1,b_1)}$ and $C_{(a_2,b_2)}$. The case where it fills the unbounded interstice, either to the left or right, is similar.

By Definition 2.11, we have $a = a_1 + a_2$ and $b = b_1 + b_2$. By Corollary 3.9 then, we get

$$\begin{aligned}\sqrt{\operatorname{curv}(C_{(a,b)})} &= \sqrt{\operatorname{curv}(C_{(a_1,b_1)})} + \sqrt{\operatorname{curv}(C_{(a_2,b_2)})}\\ &= a_1\alpha + b_1 + a_2\alpha + b_2\\ &= a\alpha + b,\end{aligned}$$

where the second equality holds by induction. □

The following is an immediate consequence of Lemma 3.13:

Corollary 3.14. *If $\alpha \notin \mathbb{Q}$, then no two circles in $\mathcal{P}_\alpha$ have the same curvature.*

Proof. Suppose $\operatorname{curv}(C_{(a,b)}) = \operatorname{curv}(C_{(a',b')})$. Then, by Lemma 3.13, we have $a\alpha + b = a'\alpha + b'$. Since α is not rational, the only way for this equation to hold is for $a = a'$ and $b = b'$ which, by the uniqueness statement in Proposition 2.21, implies that $C_{(a,b)} = C_{(a',b')}$. □

We are now in a position to prove our first main result, identifying the set of similarities between the two packings $\mathcal{P}_\alpha$ and $\mathcal{P}_\beta$ when α, β are positive real numbers. This is Theorem 1.1 from the Introduction. We start by showing how to associate a matrix to such a similarity.

Definition 3.15. Fix $\alpha, \beta > 0$ and let Φ be a similarity of $\mathbb{R}^2$ such that $\Phi \cdot \mathcal{P}_\beta = \mathcal{P}_\alpha$. Also assume that $\Phi(L) = L$, where L is the x-axis, that is, the chosen line in each packing. Then Φ takes mutually tangent circles in $\mathcal{P}_{\beta,L}$ to mutually tangent circles in $\mathcal{P}_{\alpha,L}$. In particular, we have

$$\Phi(C^{\beta}_{(1,0)}) = C^{\alpha}_{(a,b)}, \quad \Phi(C^{\beta}_{(0,1)}) = C^{\alpha}_{(c,d)}$$

for some integers a, b, c, d. By Lemma 2.13,

$$\begin{bmatrix} a & b \\ c & d \end{bmatrix}$$

is an integer matrix of determinant ± 1. The determinant is $+1$ if $C^{\alpha}_{(a,b)} \prec C^{\alpha}_{(c,d)}$, in which case Φ is orientation-preserving, and -1 if Φ is orientation-reversing. We denote this matrix by $\mathbf{A}(\Phi)$.

Theorem 3.16. *Let $\alpha, \beta > 0$. The construction* $\mathbf{A}$ *of Definition 3.15 determines a bijection between the set of similarities of* $\mathbb{R}^2$ *that take* $\mathcal{P}_\beta$ *to* $\mathcal{P}_\alpha$ *(and fix the* x*-axis) and the set of matrices* $\begin{bmatrix} a & b \\ c & d \end{bmatrix}$ *in* $\mathrm{PGL}_2(\mathbb{Z})$ *that satisfy*

$$\beta = \frac{a\alpha + b}{c\alpha + d}.$$

Furthermore, the restriction of $\mathbf{A}$ *to orientation-preserving similarities is a bijection onto the set of elements of* $\mathrm{PSL}_2(\mathbb{Z}) = \mathrm{SL}_2(\mathbb{Z})/\{\pm 1\}$ *with this property.*

Proof. We first show that $\mathbf{A}(\Phi)$ satisfies the condition that $\beta = \dfrac{a\alpha + b}{c\alpha + d}$. Because Φ is a similarity, there exists some $\lambda > 0$ such that $\mathrm{curv}(\Phi(C)) = \lambda\,\mathrm{curv}(C)$ for all $C \in \mathcal{P}_\beta$. Since $\mathrm{curv}(C^{\beta}_{(1,0)}) = \beta^2$ and $\mathrm{curv}(C^{\beta}_{(0,1)}) = 1$, it follows that $\mathrm{curv}(C^{\alpha}_{(a,b)}) = \lambda\beta^2$ and $\mathrm{curv}(C^{\alpha}_{(c,d)}) = \lambda$. By taking square roots and applying Lemma 3.13, we may conclude that

$$\begin{aligned} a\alpha + b &= \sqrt{\lambda}\beta \\ c\alpha + d &= \sqrt{\lambda}, \end{aligned}$$

which we may rewrite as $\dfrac{a\alpha + b}{c\alpha + d} = \beta$.

To show that $\mathbf{A}$ is injective, suppose $\mathbf{A}(\Phi) = \mathbf{A}(\Phi')$ in $\mathrm{PGL}_2(\mathbb{Z})$. Then, if $\Phi(C^{\beta}_{(1,0)}) = C^{\alpha}_{(a,b)}$ and $\Phi'(C^{\beta}_{(1,0)}) = C^{\alpha}_{(a',b')}$, we must have $(a, b) = \pm(a', b')$. But, by Proposition 2.21, only one of (a, b) and $(-a, -b)$ is the label of a circle in $\mathcal{P}_\alpha$. Therefore, in fact $(a, b) = (a', b')$ and so $\Phi(C^{\beta}_{(1,0)}) = \Phi'(C^{\beta}_{(1,0)})$. Similarly $\Phi(C^{\beta}_{(0,1)}) = \Phi'(C^{\beta}_{(0,1)})$. Since also $\Phi(L) = \Phi'(L)$, Lemma 3.2 tells us that $\Phi = \Phi'$.

Now let $\begin{bmatrix} a & b \\ c & d \end{bmatrix}$ be an element of $\mathrm{PGL}_2(\mathbb{Z})$ such that $\beta = \dfrac{a\alpha + b}{c\alpha + d}$. Because the determinant of this matrix is ± 1, Proposition 2.21 tells us that either (a, b) or $(-a, -b)$, but not both, is the label of a circle in $\mathcal{P}_{\alpha,L}$, and that the same holds for $(\pm c, \pm d)$. Furthermore, these circles are tangent. Because we can multiply the matrix by -1 and not change it in $\mathrm{PGL}_2(\mathbb{Z})$, we may assume that (a, b) is the label for a circle in $\mathcal{P}_{\alpha,L}$. Now we need to show that (c, d) is also the label of a circle in $\mathcal{P}_{\alpha,L}$. Suppose that $(-c, -d)$, rather than (c, d), is a label in $\mathcal{P}_{\alpha,L}$. Lemma 3.13 tells us that $a\alpha + b \geq 0$. This statement, along with the fact that $\dfrac{a\alpha + b}{c\alpha + d} = \beta > 0$, implies that $c\alpha + d > 0$, and therefore $(-c)\alpha + (-d) < 0$, which contradicts Lemma 3.13 applied to the circle $C^{\alpha}_{(-c,-d)}$. We may therefore conclude that $C^{\alpha}_{(a,b)}$ and $C^{\alpha}_{(c,d)}$ form a pair of tangent circles in $\mathcal{P}_{\alpha,L}$. Then, since the ratio of the curvatures of $C^{\alpha}_{(a,b)}$ and $C^{\alpha}_{(c,d)}$ is

$$\frac{(a\alpha + b)^2}{(c\alpha + d)^2} = \beta^2$$

which is the same as the ratio of the curvatures of $C^{\beta}_{(1,0)}$ and $C^{\beta}_{(0,1)}$, Lemma 3.5 determines a similarity Φ between $\mathcal{P}_\alpha$ and $\mathcal{P}_\beta$ such that

$$\mathbf{A}(\Phi) = \begin{bmatrix} a & b \\ c & d \end{bmatrix}.$$

This shows that $\mathbf{A}$ is a bijection. We have already noted that $\mathbf{A}(\Phi) \in \mathrm{PSL}_2(\mathbb{Z})$ if and only if Φ is orientation-preserving, which gives us the last part of the Theorem. ☐

We have the following corollaries.

Corollary 3.17. *$\mathcal{P}_\alpha$ is a strip packing if and only if $\alpha \in \mathbb{Q}^+$.*

Proof. First, we note that $\mathcal{P}_\alpha$ is a strip packing if and only if it is similar to the packing $\mathcal{P}_1$. By Theorem 3.16, this is true if and only if there is an integer matrix $\begin{bmatrix} a & b \\ c & d \end{bmatrix}$ of determinant ± 1 such that $\dfrac{a \cdot 1 + b}{c \cdot 1 + d} = \dfrac{a+b}{c+d} = \alpha$. Certainly, if such a matrix exists, then α is rational. Conversely, suppose $\alpha = \frac{p}{q}$ with $p, q > 0$ and $\gcd(p,q) = 1$. Let a and c be positive integers that satisfy $aq - cp = 1$, and set $b = p - a$, $d = q - c$. By construction, $\begin{vmatrix} a & b \\ c & d \end{vmatrix} = 1$ and $\dfrac{a+b}{c+d} = \dfrac{p}{q} = \alpha$. ☐

Corollary 3.18. *If $\mathcal{P}_\alpha$ is self-similar, then α is the root of a quadratic polynomial with rational coefficients.*

Proof. If $\mathcal{P}_\alpha$ is self-similar, then there is a nontrivial similarity Φ that maps $\mathcal{P}_\alpha$ to itself. By Theorem 3.16, this corresponds to a nontrivial element $\begin{bmatrix} a & b \\ c & d \end{bmatrix} \in \mathrm{PGL}_2(\mathbb{Z})$ such that $\dfrac{a\alpha + b}{c\alpha + d} = \alpha$; i.e., such that

$$c\alpha^2 + (d-a)\alpha - b = 0.$$

It is easy to check that the only way for all three coefficients to be zero is for $a = d = \pm 1$, $b = c = 0$, which contradicts the fact that the matrix $\begin{bmatrix} a & b \\ c & d \end{bmatrix}$ is not the identity in $\mathrm{PGL}_2(\mathbb{Z})$. ☐

The rest of this section is concerned with proving the converse of Corollary 3.18: if α is the root of a quadratic polynomial with rational coefficients, then $\mathcal{P}_\alpha$ is self-similar. This follows from Theorem 3.22 below, which is Theorem 1.2 of the Introduction.

Theorem 3.22 goes beyond identifying which packings are self-similar. We in fact calculate the self-similarity groups of all the packings. To do this we first show that, in the case $\beta = \alpha$, the bijection of Theorem 3.16 is a group isomorphism. This is the content of the following proposition.

Proposition 3.19. *Let $\alpha > 0$ be irrational. Then there are group isomorphisms*

$$\mathrm{Symm}(\mathcal{P}_\alpha) \cong \mathrm{Stab}(\alpha) := \left\{ \begin{bmatrix} a & b \\ c & d \end{bmatrix} \in \mathrm{PGL}_2(\mathbb{Z}) \;\middle|\; \frac{a\alpha + b}{c\alpha + d} = \alpha \right\}, \text{ and}$$
$$\mathrm{Symm}_+(\mathcal{P}_\alpha) \cong \mathrm{Stab}_+(\alpha) := \left\{ \begin{bmatrix} a & b \\ c & d \end{bmatrix} \in \mathrm{PSL}_2(\mathbb{Z}) \;\middle|\; \frac{a\alpha + b}{c\alpha + d} = \alpha \right\}.$$

Proof. Since α is not rational by assumption, Corollary 3.17 tell us that $\mathcal{P}_\alpha$ is not a strip packing, so the x-axis L is the unique line in the packing $\mathcal{P}_\alpha$. Therefore *every* self-similarity of $\mathcal{P}_\alpha$ maps L to L. If we take $\beta = \alpha$, then Theorem 3.16 states precisely that we have bijections of the above forms given by the construction $\mathbf{A}$. Now we show that when $\beta = \alpha$, these bijections are group isomorphisms.

Let Φ and Φ' be elements of $\mathrm{Symm}(\mathcal{P}_\alpha)$. Say

$$\mathbf{A}(\Phi) = \begin{bmatrix} a & b \\ c & d \end{bmatrix} \quad \text{and} \quad \mathbf{A}(\Phi') = \begin{bmatrix} a' & b' \\ c' & d' \end{bmatrix}.$$

We must show that $\mathbf{A}(\Phi \circ \Phi') = \begin{bmatrix} a & b \\ c & d \end{bmatrix} \begin{bmatrix} a' & b' \\ c' & d' \end{bmatrix} = \begin{bmatrix} aa' + bc' & ab' + bd' \\ ca' + dc' & cb' + dd' \end{bmatrix}$. By Corollary 3.14 and Lemma 3.2, it will suffice to show

$$\operatorname{curv}((\Phi \circ \Phi')(C_{(1,0)})) = ((aa' + bc')\alpha + (ab' + bd'))^2 \text{ and} \tag{3.20}$$

$$\operatorname{curv}((\Phi \circ \Phi')(C_{(0,1)})) = ((ca' + dc')\alpha + (cb' + dd'))^2. \tag{3.21}$$

In fact, it will suffice to show only that (3.21) holds: since $\Phi \circ \Phi'$ is a similarity, we must have $\operatorname{curv}((\Phi \circ \Phi')(C_{(1,0)})) = \lambda \operatorname{curv}(C_{(1,0)}) = \lambda\alpha^2$ and $\operatorname{curv}((\Phi \circ \Phi')(C_{(0,1)})) = \lambda \operatorname{curv}(C_{(0,1)}) = \lambda$ for some $\lambda > 0$; therefore, if (3.21) is satisfied, the fact that the product of the matrices is still an element of $\mathrm{Stab}(\alpha)$ will force (3.20).

The scale factor λ under the composition $\Phi \circ \Phi'$ is the product of the scale factors μ and μ' under the maps Φ and Φ' respectively. Since $\operatorname{curv}(C_{(0,1)}) = 1$, it follows that

$$\mu = \operatorname{curv}(\Phi(C_{(0,1)})) = \operatorname{curv}(C_{(c,d)}) = (c\alpha + d)^2 \text{ and}$$
$$\mu' = \operatorname{curv}(\Phi'(C_{(0,1)})) = \operatorname{curv}(C_{(c',d')}) = (c'\alpha + d')^2.$$

Therefore we will have shown that $\mathbf{A}$ is a group homomorphism if we can show that

$$(c\alpha + d)(c'\alpha + d') = (ca' + dc')\alpha + (cb' + dd').$$

Indeed,

$$\begin{aligned} \frac{a'\alpha + b'}{c'\alpha + d'} = \alpha &\implies c'\alpha^2 + (d' - a')\alpha - b' = 0 \\ &\implies cc'\alpha^2 + (cd' - ca')\alpha - cb' = 0 \\ &\implies cc'\alpha^2 + (cd' + dc' - ca' - dc')\alpha - cb' = 0 \\ &\implies cc'\alpha^2 + (cd' + dc')\alpha = (ca' + dc')\alpha + cb' \\ &\implies cc'\alpha^2 + (cd' + dc')\alpha + dd' = (ca' + dc')\alpha + (cb' + dd') \\ &\implies (c\alpha + d)(c'\alpha + d') = (ca' + dc')\alpha + (cb' + dd'). \end{aligned}$$

□

Now we can prove the main result of this section.

Theorem 3.22. *Let $\alpha > 0$. Then:*

(i) *If $\alpha \in \mathbb{Q}$, then $\mathcal{P}_\alpha$ is a strip packing and*

$$\mathrm{Symm}(\mathcal{P}_\alpha) \cong D_\infty \times \mathbb{Z}/2\mathbb{Z}$$

with subgroup

$$\mathrm{Symm}_+(\mathcal{P}_\alpha) \cong D_\infty,$$

where D_∞ denotes the infinite dihedral group.

(ii) *If α is quadratic over $\mathbb{Q}$, then*

$$\mathrm{Symm}(\mathcal{P}_\alpha) \cong \mathbb{Z}.$$

Let D denote the discriminant of the primitive integral polynomial with root α. Then the subgroup $\mathrm{Symm}_+(\mathcal{P}_\alpha)$ *is:*

- *equal to* $\mathrm{Symm}(\mathcal{P}_\alpha)$ *if the Pell equation $x^2 - Dy^2 = -4$ has no integral solution for (x, y);*

- *the index* 2 *subgroup of* $\mathrm{Symm}(\mathcal{P}_\alpha)$ *if* $x^2 - Dy^2 = -4$ *does have an integral solution.*

(iii) *Otherwise*

$$\mathrm{Symm}_+(\mathcal{P}_\alpha) = \mathrm{Symm}(\mathcal{P}_\alpha) = 1.$$

Proof. First of all, it follows immediately from Corollary 3.18 that, for any $\alpha > 0$ that is neither rational nor quadratic, $\mathcal{P}_\alpha$ is not self-similar. Therefore $\mathrm{Symm}(\mathcal{P}_\alpha) = \mathrm{Symm}_+(\mathcal{P}_\alpha) = 1$ for all such α.

Now suppose that α is rational. Then Corollary 3.17 tells us that $\mathcal{P}_\alpha$ is a strip packing. In this case, the full self-similarity group is generated by a translation 'along' the strip, a reflection in a line perpendicular to the strip, and a reflection that interchanges the two lines. The resulting group is isomorphic to $D_\infty \times \mathbb{Z}/2\mathbb{Z}$ where D_∞ is the infinite dihedral group. The subgroup of orientation-preserving self-similarities of $\mathcal{P}_\alpha$ is generated by the translation and the rotation given by combining the two reflections. This subgroup is isomorphic to D_∞.

The main focus of our work is the case where α is of degree precisely 2 over $\mathbb{Q}$. We have already shown in Proposition 3.19 that

$$\mathrm{Symm}(\mathcal{P}_\alpha) \cong \left\{ \begin{bmatrix} a & b \\ c & d \end{bmatrix} \in \mathrm{PGL}_2(\mathbb{Z}) \;\middle|\; \frac{a\alpha + b}{c\alpha + d} = \alpha \right\}$$

and that

$$\mathrm{Symm}_+(\mathcal{P}_\alpha) \cong \left\{ \begin{bmatrix} a & b \\ c & d \end{bmatrix} \in \mathrm{PSL}_2(\mathbb{Z}) \;\middle|\; \frac{a\alpha + b}{c\alpha + d} = \alpha \right\},$$

so it suffices to calculate these stabilizer groups.

It turns out that the elements of $\mathrm{PGL}_2(\mathbb{Z})$ that fix α are closely related to the solutions to the Pell equations

$$x^2 - Dy^2 = \pm 4, \tag{3.23}$$

where $D = q^2 - 4pr$ is the discriminant of the primitive integer polynomial $f(x) = px^2 + qx + r$ satisfied by α, with $p > 0$.

Define the sets $\mathcal{G}$ and $\mathcal{G}_+$ as follows:

$$\mathcal{G} := \left\{ \frac{x + y\sqrt{D}}{2} \;\middle|\; x, y \in \mathbb{Z},\ x^2 - Dy^2 = \pm 4 \right\}$$

and

$$\mathcal{G}_+ := \left\{ \frac{x + y\sqrt{D}}{2} \;\middle|\; x, y \in \mathbb{Z},\ x^2 - Dy^2 = 4 \right\},$$

where D is as in the previous paragraph. One can easily check that $\mathcal{G}$ is a group under multiplication with subgroup $\mathcal{G}_+$. Furthermore, one can show (see [8, Theorem 1.9], for example) that

$$\mathcal{G} \cong \mathbb{Z} \times \{\pm 1\}. \tag{3.24}$$

Our proof of Theorem 3.22 is given by relating the stabilizer of α in $\mathrm{PGL}_2(\mathbb{Z})$ with the group $\mathcal{G}$.

Recall that the integers p, q, r are the coefficients of the primitive integer polynomial satisfied by α and that $D = q^2 - 4pr$. We first construct a group homomorphism

$$\begin{aligned} \Gamma : \mathcal{G} &\to \left\{ \begin{bmatrix} a & b \\ c & d \end{bmatrix} \in \mathrm{PGL}_2(\mathbb{Z}) \;\middle|\; \frac{a\alpha + b}{c\alpha + d} = \alpha \right\} \\ \frac{x + y\sqrt{D}}{2} &\mapsto \begin{bmatrix} \frac{x - yq}{2} & -yr \\ yp & \frac{x + yq}{2} \end{bmatrix} \end{aligned}$$

The matrices in the image of Γ consist of integer entries because:

$$\begin{aligned} x \pm yq &\equiv x^2 - y^2q^2 \pmod 2 \\ &\equiv x^2 - y^2(q^2 - 4pr) \pmod 2 \\ &= x^2 - Dy^2 = \pm 4 \equiv 0 \pmod 2. \end{aligned}$$

These matrices stabilize α because:

$$\begin{aligned} p\alpha^2 + q\alpha + r = 0 &\implies yp\alpha^2 + yq\alpha + yr = 0 \\ &\implies yp\alpha^2 + \left(\frac{x+yq}{2} - \frac{x-yq}{2}\right)\alpha + yr = 0 \\ &\implies \frac{\frac{x-yq}{2}\alpha - yr}{yp\alpha + \frac{x+yq}{2}} = \alpha, \end{aligned}$$

and are invertible because

$$\begin{aligned} \begin{vmatrix} \frac{x-yq}{2} & -yr \\ yp & \frac{x+yq}{2} \end{vmatrix} &= \frac{x^2 - y^2q^2}{4} + y^2pr \\ &= \frac{1}{4}(x^2 - Dy^2) \\ &= \pm 1, \end{aligned}$$

where the sign is positive if and only if (x, y) satisfies $x^2 - Dy^2 = +4$. To see that Γ is a group homomorphism, we check:

$$\begin{aligned} \Gamma\left(\frac{x + y\sqrt{D}}{2} \cdot \frac{x' + y'\sqrt{D}}{2}\right) &= \Gamma\left(\frac{\frac{xx'+Dyy'}{2} + \frac{xy'+x'y}{2}\sqrt{D}}{2}\right) \\ &= \begin{bmatrix} \frac{xx'+yy'(q^2-4pr)-xy'q-x'yq}{4} & \frac{-xy'r-x'yr+yy'qr-yy'qr}{2} \\ \frac{x'yp+xy'p+yy'pq-yy'pq}{2} & \frac{xx'+yy'(q^2-4pr)+xy'q+x'yq}{4} \end{bmatrix} \\ &= \begin{bmatrix} \frac{x-yq}{2} & -yr \\ yp & \frac{x+yq}{2} \end{bmatrix} \begin{bmatrix} \frac{x'-y'q}{2} & -y'r \\ y'p & \frac{x'+y'q}{2} \end{bmatrix} \\ &= \Gamma\left(\frac{x + y\sqrt{D}}{2}\right)\Gamma\left(\frac{x' + y'\sqrt{D}}{2}\right). \end{aligned}$$

Combining the map Γ with the isomorphism of Proposition 3.19 we have now shown how to construct, for each solution to (3.23), a self-similarity of $\mathcal{P}_\alpha$. To prove our Theorem, we calculate the kernel and image of the homomorphism Γ. First, we show that Γ is surjective, which implies that every self-similarity of $\mathcal{P}_\alpha$ arises from a solution to (3.23) in the manner described above. So suppose we are given a matrix $A = \begin{bmatrix} a & b \\ c & d \end{bmatrix}$ that stabilizes α. In particular, it follows that

$$c\alpha^2 + (d - a)\alpha - b = 0$$

This polynomial is therefore an integer multiple of the primitive polynomial $px^2 + qx + r$ with root α. That is, there exists $m \in \mathbb{Z}$ such that

$$\begin{aligned} c &= mp \\ d - a &= mq \\ -b &= mr. \end{aligned}$$

Now set

$$\begin{aligned} x &= a + d \\ y &= m. \end{aligned}$$

We clearly have $x, y \in \mathbb{Z}$ and

$$\begin{aligned} x^2 - Dy^2 &= (a+d)^2 - (q^2 - 4pr)m^2 \\ &= (a+d)^2 - (a-d)^2 - 4bc \\ &= 4(ad - bc) = \pm 4. \end{aligned}$$

so $\frac{x+y\sqrt{D}}{2} \in \mathcal{G}$. (Moreover, this is in $\mathcal{G}_+$ if and only if $A \in \mathrm{PSL}_2(\mathbb{Z})$.) It is easy to check that $\Gamma(\frac{x+y\sqrt{D}}{2}) = A$ as required.

Finally, $\frac{x+y\sqrt{D}}{2}$ is in the kernel of Γ if and only if

$$x - yq = x + yq = \pm 2\ ,\ \ yr = yp = 0.$$

Since p cannot be zero ($\alpha \notin \mathbb{Q}$), it follows that $y = 0$, and therefore $x = \pm 2$. In other words,

$$\ker(\Gamma) = \{\pm 1\}.$$

Putting together our various isomorphisms and using (3.24), we now have

$$\mathrm{Symm}(\mathcal{P}_\alpha) \cong \mathcal{G}/\{\pm 1\} \cong \mathbb{Z}.$$

We have also seen that the orientation-preserving self-similarities correspond under this isomorphism to the subgroup $\mathcal{G}_+/\{\pm 1\}$. There are two possibilities here. One is that the generator for $\mathcal{G}$ is in $\mathcal{G}_+$. In this case the groups are equal and all the self-similarities of $\mathcal{P}_\alpha$ are orientation-preserving. This happens when there are no integer solutions to the equation

$$x^2 - Dy^2 = -4.$$

The other possibility is that the generator $z = \frac{x_0+y_0\sqrt{D}}{2}$ for $\mathcal{G}$ is not in $\mathcal{G}_+$. But then, however, z^2 is in $\mathcal{G}_+$ and so $\mathcal{G}_+$ is an index 2 subgroup of $\mathcal{G}$. In this case, $\mathrm{Symm}_+(\mathcal{P}_\alpha)$ is an index 2 subgroup in $\mathrm{Symm}(\mathcal{P}_\alpha)$ as claimed. ☐

Corollary 3.25. *The half-plane packing $\mathcal{P}_\alpha$ is self-similar if and only if α is rational or quadratic over $\mathbb{Q}$. In the quadratic case, $\mathcal{P}_\alpha$ is self-similar via an orientation-reversing self-similarity if and only if the equation $x^2 - Dy^2 = -4$ has an integral solution (x, y), where D is the discriminant of the primitive integral polynomial with root α.*

4. Half-plane Packings and Continued Fractions

In this section, our main goal is to describe how the continued fraction of a positive real number α manifests itself geometrically in the half-plane packing $\mathcal{P}_\alpha$. Recall that the standard continued fraction expansion of a positive real number α is a representation of the form

$$\alpha = a_0 + \cfrac{1}{a_1 + \cfrac{1}{a_2 + \cfrac{1}{\ddots}}}$$

for some integers a_k, with $a_0 \geq 0$ and $a_k > 0$ for all $k > 0$. It is also standard (and more practical typographically) to express this expansion simply as $\alpha = [a_0, a_1, a_2, \ldots]$.

To begin, we recall the algorithm for computing the continued fraction expansion of a positive real number α. The continued fraction expansion is computed by successive iterations of the following algorithm, which we refer to as the **continued fraction algorithm**. The input for the algorithm is the number $\alpha_0 = \alpha$. Each step of the algorithm takes α_n and determines α_{n+1}.

(A) If $\alpha_n \geq 1$, let $\alpha_{n+1} = \alpha_n - 1$.
(B) If $0 < \alpha_n < 1$, let $\alpha_{n+1} = \frac{1}{\alpha_n}$.
(C) If $\alpha_n = 0$, halt.

Recording the sequence of steps obtained when applying this algorithm to a positive real number α we get something like

$$ABAABAAC.$$

The positive integer a_k from the continued fraction expansion corresponds precisely to the length of the $(k+1)$th string of consecutive A's. For example, the above sequence represents the application of the continued fraction algorithm to $\alpha_0 = \frac{7}{5}$. The resulting continued fraction expansion is

$$[1,2,2] = 1 + \cfrac{1}{2 + \cfrac{1}{2}} = \frac{7}{5}.$$

The sequence (α_n) in this case is:

$$\frac{7}{5} \xrightarrow{\text{(A)}} \frac{2}{5} \xrightarrow{\text{(B)}} \frac{5}{2} \xrightarrow{\text{(A)}} \frac{3}{2} \xrightarrow{\text{(A)}} \frac{1}{2} \xrightarrow{\text{(B)}} 2 \xrightarrow{\text{(A)}} 1 \xrightarrow{\text{(A)}} 0. \tag{4.1}$$

Notice that the continued fraction expansion for α_n is the same as that for α, but with an 'initial segment' removed. For example, if $\alpha_0 = [2,3,4,5,6,7]$, then we have

$$\begin{aligned}
\alpha_1 &= [1,3,4,5,6,7] \\
\alpha_2 &= [0,3,4,5,6,7] \\
\alpha_3 &= [3,4,5,6,7] \\
\alpha_4 &= [2,4,5,6,7] \\
\alpha_5 &= [1,4,5,6,7] \\
\alpha_6 &= [0,4,5,6,7] \\
\alpha_7 &= [4,5,6,7]
\end{aligned}$$

and so on.

Turning now back to Apollonian circle packings, we define a **circle replacement algorithm**. The input of this algorithm is an ordered pair (X_0, Y_0), where X_0 and Y_0 are tangent circles in a half-plane packing $\mathcal{P}$ that are also tangent to a chosen line $L \in \mathcal{P}$. We also require that Y_0 is not itself a line. At the $(n+1)$th step of the algorithm, we replace the pair (X_n, Y_n) with a new pair of circles (X_{n+1}, Y_{n+1}):

(A) If $\operatorname{curv}(X_n) \geq \operatorname{curv}(Y_n)$, take $Y_{n+1} = Y_n$ and take X_{n+1} to be the circle that fills the unbounded interstice for X_n and Y_n (in the sense of Definition 2.8). Note that Corollary 3.9 implies that

$$\sqrt{\operatorname{curv}(X_{n+1})} = \sqrt{\operatorname{curv}(X_n)} - \sqrt{\operatorname{curv}(Y_n)}.$$

(B) If $0 < \operatorname{curv}(X_n) < \operatorname{curv}(Y_n)$, take $X_{n+1} = Y_n$ and $Y_{n+1} = X_n$.
(C) If $\operatorname{curv}(X_n) = 0$, halt.

As with the continued fraction algorithm, we are interested in the sequence of steps involved when the algorithm is performed to a given starting pair of circles. (For example, we might obtain the sequence $AABABAAAABAAC$.) Our main observation is then the following.

Lemma 4.2. *Let α be a positive real number, and let (X_0, Y_0) be the two circles used to construct the half-plane packing $\mathcal{P}_\alpha$: X_0 and Y_0 are tangent to each other and to the x-axis L, and $\operatorname{curv}(X_0) = \alpha^2$, $\operatorname{curv}(Y_0) = 1$. Then the sequence of steps (A, B, or C) performed in applying the continued fraction algorithm to α is the same as the sequence of steps performed in applying the circle replacement algorithm to (X_0, Y_0). Moreover, we have*

$$\alpha_n = \frac{\sqrt{\operatorname{curv}(X_n)}}{\sqrt{\operatorname{curv}(Y_n)}}$$

for all $n \geq 0$.

Proof. The proof is by induction on n. For $n = 0$, this is the claim

$$\alpha = \frac{\sqrt{\operatorname{curv}(X_0)}}{\sqrt{\operatorname{curv}(Y_0)}}$$

which is true by the choice of X_0 and Y_0.

Suppose that the claim holds for α_n and (X_n, Y_n). Then $\alpha_n \geq 1$ if and only if $\operatorname{curv}(X_n) \geq \operatorname{curv}(Y_n)$ and $\alpha_n = 0$ if and only if $\operatorname{curv}(X_n) = 0$. This tells us that the next step (A, B, or C) will be the same for both algorithms. So it remains only to verify that the formula still holds for α_{n+1} and (X_{n+1}, Y_{n+1}).

Suppose that $\alpha_n \geq 1$. Then we have $\alpha_{n+1} = \alpha_n - 1$, so it is sufficient to show that

$$\frac{\sqrt{\operatorname{curv}(X_{n+1})}}{\sqrt{\operatorname{curv}(Y_{n+1})}} = \frac{\sqrt{\operatorname{curv}(X_n)}}{\sqrt{\operatorname{curv}(Y_n)}} - 1.$$

We have $Y_{n+1} = Y_n$, so it is enough to show that

$$\sqrt{\operatorname{curv}(X_{n+1})} = \sqrt{\operatorname{curv}(X_n)} - \sqrt{\operatorname{curv}(Y_n)}$$

which follows from Corollary 3.9 as mentioned above.

Finally, suppose that $0 < \alpha_n < 1$. Then

$$\alpha_{n+1} = \frac{1}{\alpha_n} = \frac{\sqrt{\operatorname{curv}(Y_n)}}{\sqrt{\operatorname{curv}(X_n)}} = \frac{\sqrt{\operatorname{curv}(X_{n+1})}}{\sqrt{\operatorname{curv}(Y_{n+1})}}.$$

□

Figure 10 shows the circle replacement algorithm applied to the packing $\mathcal{P}_\alpha$ for $\alpha = \frac{7}{5}$. The circles are labeled by the square roots of their curvatures. Compare this to (4.1) as an illustration of Lemma 4.2.

Recall that the continued fraction expansion of a real number α determines a sequence of rational numbers $\frac{p_n}{q_n}$ that converge to α. These are the *convergents* of α and are given by truncating the continued fraction expansion of α. Thus if

$$\alpha = [a_0, a_1, a_2, \dots]$$

then set

$$\frac{p_n}{q_n} := [a_0, a_1, \dots, a_n],$$

where p_n and q_n are nonnegative coprime integers. They satisfy the recurrence equations

$$p_n = p_{n-2} + a_n p_{n-1}, \quad q_n = q_{n-2} + a_n q_{n-1}. \tag{4.3}$$

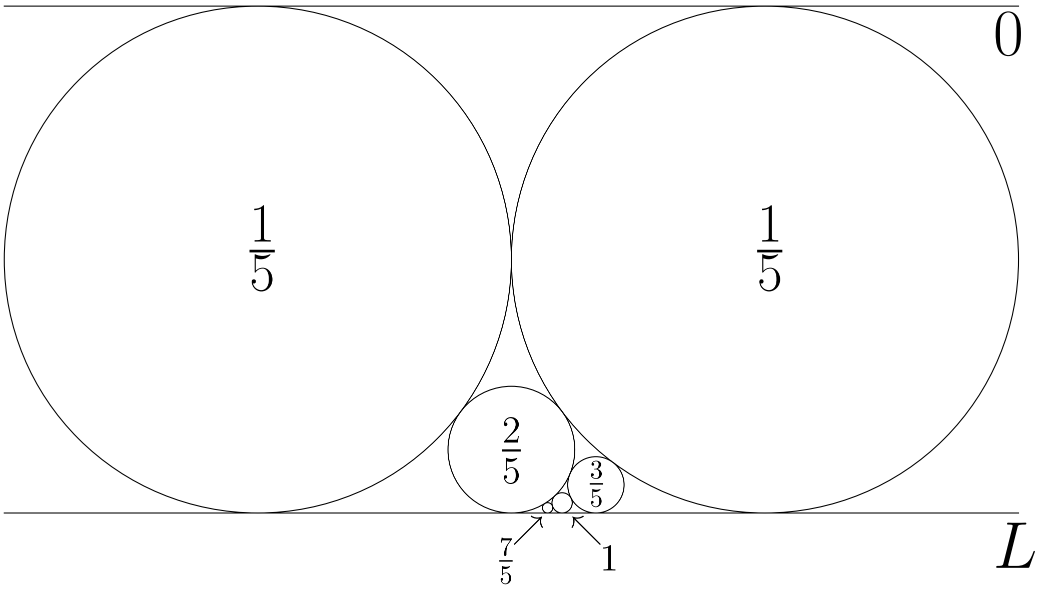

FIGURE 10. The circle replacement algorithm applied to the packing $\mathcal{P}_{\frac{7}{5}}$. The circles are labeled by the square roots of their curvatures.

We now observe that the convergents of α appear in the labels (in the sense of §2) of the circles in the circle replacement algorithm applied to the packing $\mathcal{P}_\alpha$.

Lemma 4.4. *Let α be a positive real number. Let $(X_0, Y_0) = (C_{(1,0)}, C_{(0,1)})$ be the generating circles for the packing $\mathcal{P}_\alpha$. The sequence of* distinct *circles in the sequence (Y_j) defined by the circle replacement algorithm is*

$$C_{(0,1)}, C_{(q_0,-p_0)}, C_{(-q_1,p_1)}, C_{(q_2,-p_2)}, C_{(-q_3,p_3)}, \ldots.$$

In particular, if $\alpha \notin \mathbb{Q}$, then

$$\lim_{j\to\infty} \operatorname{curv}(Y_j) = 0.$$

Proof. We have $Y_0 = C_{(0,1)}$ and $X_0 = C_{(1,0)}$. The first new Y_j will appear after the first application of step (B) of the algorithm, that is, after $a_0 + 1$ steps. At this point we have

$$Y_{a_0+1} = X_{a_0} = C_{(1,-a_0)} = C_{(q_0,-p_0)}$$

and

$$X_{a_0+1} = Y_{a_0} = C_{(0,1)}.$$

Now suppose, inductively, that immediately after the nth application of step (B) we have

$$Y_N = C_{(q_{n-1},-p_{n-1})}, X_N = C_{(-q_{n-2},p_{n-2})}.$$

Running the algorithm until after the next application of (B), that is $a_n + 1$ times, we have

$$Y_{N+a_n+1} = X_{N+a_n} = C_{(-q_{n-2}-a_nq_{n-1},p_{n-2}+a_np_{n-1})} = C_{(-q_n,p_n)}$$

and

$$X_{N+a_n+1} = Y_{N+a_n} = Y_N = C_{(q_{n-1},-p_{n-1})}.$$

The first claim now follows by induction on n.

By Lemma 3.13, we have

$$\operatorname{curv}(C_{(\pm q_n,\mp p_n)}) = (q_n\alpha - p_n)^2 = q_n^2\left(\alpha - \frac{p_n}{q_n}\right)^2.$$

A basic fact about the convergents for continued fractions [7, Theorem 171] is that

$$\left|\alpha - \frac{p_n}{q_n}\right| < \frac{1}{q_n^2}$$

so

$$\operatorname{curv}(C_{(\pm q_n, \mp p_n)}) < \frac{1}{q_n^2}.$$

It follows from (4.3) that $q_n \to \infty$ as $n \to \infty$, so the curvatures of the Y_j tend to zero. ☐

Our goal is now to use this relationship between the circle replacement and continued fraction algorithms to give new criteria for two packings to be similar, and a new way to understand the self-similarities of a given packing, both in terms of continued fraction expansions. To do this we have to know that the circle replacement algorithm involves 'enough' of the circles in the packing to be able to detect any similarity between two packings. The following lemma is key to this.

Lemma 4.5. *Let $\mathcal{P}$ be a half-plane packing (but not a strip packing). Fix an ordered pair of circles (X_0, Y_0) as in the definition of the circle replacement algorithm. Let X and Y be any pair of tangent circles in $\mathcal{P}_L$ such that X_0 and Y_0 are contained in the bounded interstice formed by X and Y. Then one of the pairs (X, Y) and (Y, X) appears as (X_n, Y_n) in the application of the circle replacement algorithm to (X_0, Y_0).*

Proof. The second part of Lemma 4.4 implies that the radii of the circles X_n and Y_n increase without bound as n tends to infinity. Therefore there is some smallest integer N such that X_N does not lie in the bounded interstice formed by X and Y. Since X_N is tangent to X_{N-1}, the remark following Definition 2.8 implies that X_N also cannot lie in the unbounded interstice for X and Y. Hence X_N is equal to one of X, Y, say Y without loss of generality.

Now Y_N is either equal to X_n for some $n < N$, or is equal to Y_0. Either way, Y_N is in the bounded interstice formed by X and Y and so, in particular, is smaller than X_N. This tells us that the next step in the circle replacement algorithm is (B); i.e., $X_{N+1} = Y_N$ and $Y_{N+1} = X_N = Y$. Since Y_{N+1} is larger than X_{N+1}, we next repeat step (A) until X_{N+K} is larger than $Y_{N+K} = Y_{N+1}$ for some $K \geq 0$. Then X_{N+K} is not in the bounded interstice formed by X and Y, but X_{N+1} is. This means we can find a smallest M with $N + 1 < M \leq N + K$ such that X_M is not in the bounded interstice formed by X and Y. Since it is tangent to X_{M-1}, this circle X_M also cannot be in the unbounded interstice, so must be one of X and Y. But it is not Y since Y_M is. Therefore we have $X_M = X$ and $Y_M = Y$ which completes the proof. ☐

We can now relate properties of the continued fraction expansion of a positive real number α to geometric properties of the half-plane circle packing $\mathcal{P}_\alpha$.

4.1. Strip packings.

We already saw in Corollary 3.17 that $\mathcal{P}_\alpha$ is the strip packing if and only if $\alpha \in \mathbb{Q}$. This is now reflected in the fact that the continued fraction expansion for α halts if and only if $\alpha \in \mathbb{Q}$. We can see from Lemma 4.2 that the continued fraction expansion of α halts exactly when the corresponding circle replacement algorithm produces a circle of curvature 0, that is, a straight line. This is illustrated in the example of $\alpha = \frac{7}{5}$ displayed above.

4.2. Similar packings.

We can determine whether the packings $\mathcal{P}_\alpha$ and $\mathcal{P}_\beta$ are similar by examining the tails of the continued fractions of α and β.

Definition 4.6. Let us say that α and β have **eventually equal** continued fraction expansions if there is some $k, N \in \mathbb{Z}$ such that $a_n = b_{n+k}$ for all $n \geq N$ (where $[a_i]$ is the continued fraction expansion of α and $[b_j]$ is the continued fraction expansion of β). For example, $[1, 2, 3, 4, 4, 4, 4, \dots]$

and $[2, 7, 4, 4, 4, 4, \dots]$ are eventually equal. For the purposes of Theorem 4.7, we say that two finite continued fraction expansions are eventually equal since both expansions terminate.

It is easy to see that α and β have eventually equal continued fraction expansions if and only if there exist $m, n \geq 0$ such that $\alpha_n = \beta_m$ (where these are the sequences obtained by applying the continued fraction algorithm to α and β). This observation yields the following classification of half-plane packings up to similarity.

Theorem 4.7. *For positive real numbers α, β, the circle packings $\mathcal{P}_\alpha$ and $\mathcal{P}_\beta$ are similar if and only if α and β have eventually equal continued fraction expansions.*

Proof. A quick proof of this result follows by identifying each of the conditions in the statement with the condition that there exist $p, q, r, s \in \mathbb{Z}$ with $ps - qr = \pm 1$ and $\frac{p\alpha+q}{r\alpha+s} = \beta$. For the circle packings, this is Lemma 3.16; for the continued fractions, it is [7, Theorem 175]. However, we give a more interesting proof arising from the direct comparison between the continued fraction and circle replacement algorithms.

First, note that α and β have finite continued fraction expansions if and only if α and β are rational, which is equivalent by Corollary 3.17 to $\mathcal{P}_\alpha$ and $\mathcal{P}_\beta$ both being strip packings, which are similar. We may therefore assume that α and β are irrational.

Suppose α and β have eventually equal continued fraction expansions. Then $\alpha_n = \beta_m$ for some m, n. This means that the ratio of the curvatures of a pair of tangent circles in $\mathcal{P}_\alpha$, both tangent to L, is equal to the ratio of the curvatures of a pair of tangent circles in $\mathcal{P}_\beta$, both tangent to L. It follows by Lemma 3.5 that there is a similarity between $\mathcal{P}_\alpha$ and $\mathcal{P}_\beta$.

To prove the converse, suppose $\mathcal{P}_\alpha$ and $\mathcal{P}_\beta$ are similar. Then there is a pair of circles (X_0', Y_0') in $\mathcal{P}_\alpha$ whose ratio of curvatures is equal to β^2, in addition to the original pair of circles (X_0, Y_0) in $\mathcal{P}_\alpha$ whose ratio of curvatures is α^2. The key step is the following claim: if we apply the circle replacement algorithm to each of these pairs of circles, they will eventually coincide; that is, there is some pair of circles (X, Y) in $\mathcal{P}_\alpha$ that appears both as (X_M, Y_M) and (X_N', Y_N') for some $M, N \in \mathbb{N}$. Note that the circle replacement algorithm only sees *ratios* of curvatures and not the curvatures themselves, so the circle replacement algorithm will generate the same numerical data for (X', Y') as it would for the corresponding pair of circles in the packing $\mathcal{P}_\beta$. From this claim, it follows that α and β have eventually equal continued fraction expansions since

$$\alpha_M = \frac{\sqrt{\operatorname{curv}(X_M)}}{\sqrt{\operatorname{curv}(Y_M)}} = \frac{\sqrt{\operatorname{curv}(X)}}{\sqrt{\operatorname{curv}(Y)}} = \frac{\sqrt{\operatorname{curv}(X_N')}}{\sqrt{\operatorname{curv}(Y_N')}} = \beta_N.$$

Now let us prove that the circle replacement algorithms corresponding to α and β eventually coincide in the sense described in the previous paragraph. Let (X_n, Y_n) be the pairs obtained from applying the circle replacement algorithm to (X_0, Y_0), as defined in the previous paragraph. We suppose without loss of generality that $X_0 \prec Y_0$ and $X_0' \prec Y_0'$. If one of the pairs contained the other in its bounded interstice, say $X_0 \preceq X_0' \prec Y_0' \preceq Y_0$, then by Lemma 4.5 there would be M such that $\{X_M, Y_M\} = \{X_0', Y_0'\}$. Suppose instead that

$$X_0 \prec Y_0 \preceq X_0' \prec Y_0'.$$

Suppose that there is no pair (X_n, Y_n) that contains (X_0', Y_0') in its bounded interstice. Then, for each n, one of the circles X_n and Y_n, has its point of tangency with the x-axis between the corresponding tangency points of Y_0 and X_0'. By the second part of Lemma 4.4, this means that there are arbitrarily large circles, all disjoint, with tangency points in this fixed interval. A little geometry shows that if two disjoint circles of radii R and R' are tangent to the x-axis, then their points of tangency are at least $2\sqrt{RR'}$ apart. This gives us a contradiction and so we deduce that there is M such that X_M and Y_M contain both X_0' and Y_0' in their bounded interstice. But then by Lemma 4.5, the pair $\{X_M, Y_M\}$ is equal to $\{X_N', Y_N'\}$ for some N. To complete the proof that the algorithms eventually coincide, we need to show that we can choose M and N such that $X_M = X_N'$ and $Y_M = Y_N'$.

Suppose instead that $X_M = Y_N'$ and $Y_M = X_N'$. We may assume that X_M is smaller than Y_M (otherwise, apply one more replacement to (X_M, Y_M) to replace X_M with Y_{M+1} and Y_M with X_{M+1}). It follows that Y_N' is smaller than X_N', so that the next step of the algorithm, step (B), will set $X_{N+1}' = Y_N'$ and $Y_{N+1}' = X_N'$, and we therefore have $X_M = X_{N+1}'$ and $Y_M = Y_{N+1}'$. ☐

4.3. **Self-similar packings.** We also have already seen (Theorem 3.22) that $\mathcal{P}_\alpha$ is self-similar (but not the strip packing) if and only if α is quadratic over $\mathbb{Q}$. It is a well-known fact that an irrational number α is quadratic over $\mathbb{Q}$ if and only if its continued fraction expansion is infinite and periodic; i.e., if and only if

$$\alpha = [a_0, \dots, a_{n-1}, c_0, \dots, c_{m-1}, c_0, \dots, c_{m-1}, \dots] = [a_0, \dots, a_{n-1}, \overline{c_0, \dots, c_{m-1}}].$$

This fact, along with Theorem 4.7, yields the following classification of self-similar half-plane packings (which are not strip packings):

Theorem 4.8. *The similarity classes of self-similar half-plane (non-strip) circle packings correspond bijectively to finite non-repeating sequences of positive integers, up to cyclic permutation. (A sequence is non-repeating if it is not equal to the concatenation of multiple copies of the same smaller sequence.)*

Proof. $\mathcal{P}_\alpha$ is self-similar if and only if the continued fraction expansion for α is periodic, as we mentioned above. We identify the similarity class of $\mathcal{P}_\alpha$ with the minimal periodic part of this expansion. For example, if $\alpha = \sqrt{2} = [1, 2, 2, 2, \dots]$ then we identify $[\mathcal{P}_\alpha]$ with the one term sequence (2). If $\alpha = \sqrt{3} = [1, 1, 2, 1, 2, 1, 2, \dots]$, we identify $[\mathcal{P}_\alpha]$ with (1, 2), or equivalently, (2, 1). Conversely, the finite non-repeating sequence $(a_0, \dots, a_{n-1})$ represents the quadratic number $\alpha = [\overline{a_0, \dots, a_{n-1}}]$, so that every such sequence represents a similarity class of self-similar half-plane packings (namely, the class including $\mathcal{P}_\alpha$). That each sequence represents exactly one similarity class follows from Theorem 4.7, since two periodic continued fractions expansions are eventually equal if and only if they have the same periodic part up to a cyclic permutation. ☐

Examples. Based on the classification in Theorem 4.8 we can give examples of the simplest self-similar half-plane packings. In some sense, the simplest such packing is given by $\alpha = [1, 1, \dots] = \frac{1+\sqrt{5}}{2}$. From the perspective of Theorem 4.8, this is represented by the singleton sequence (1). The corresponding circle packing $\mathcal{P}_\alpha$ has a self-similarity constructed from a single circle replacement. This is displayed in Figure 11.

The next simplest example is $\alpha = [2, 2, \dots] = 1 + \sqrt{2}$, represented by the singleton sequence (2). The corresponding circle packing $\mathcal{P}_\alpha$ has a self-similarity obtained by doing two circle replacements. This appears in Figure 12.

There are two different self-similar packings for which a self-similarity involves three circle replacements. Corresponding to the sequence (3), we have $\alpha = [3, 3, \dots] = \frac{3+\sqrt{13}}{2}$. The packing $\mathcal{P}_\alpha$ is shown in Figure 13. Corresponding to the sequence (1,2) we have $\alpha = [1, 2, 1, 2, \dots] = \frac{1+\sqrt{3}}{2}$. This packing is shown in Figure 14.

We have seen that the self-similarity of the packing $\mathcal{P}_\alpha$ is reflected in the continued fraction expansion of α. In fact, it turns out that every self-similarity of $\mathcal{P}_\alpha$ comes about from the periodicity of the continued fraction expansion. In particular, the self-similarity groups of the packings in Figures

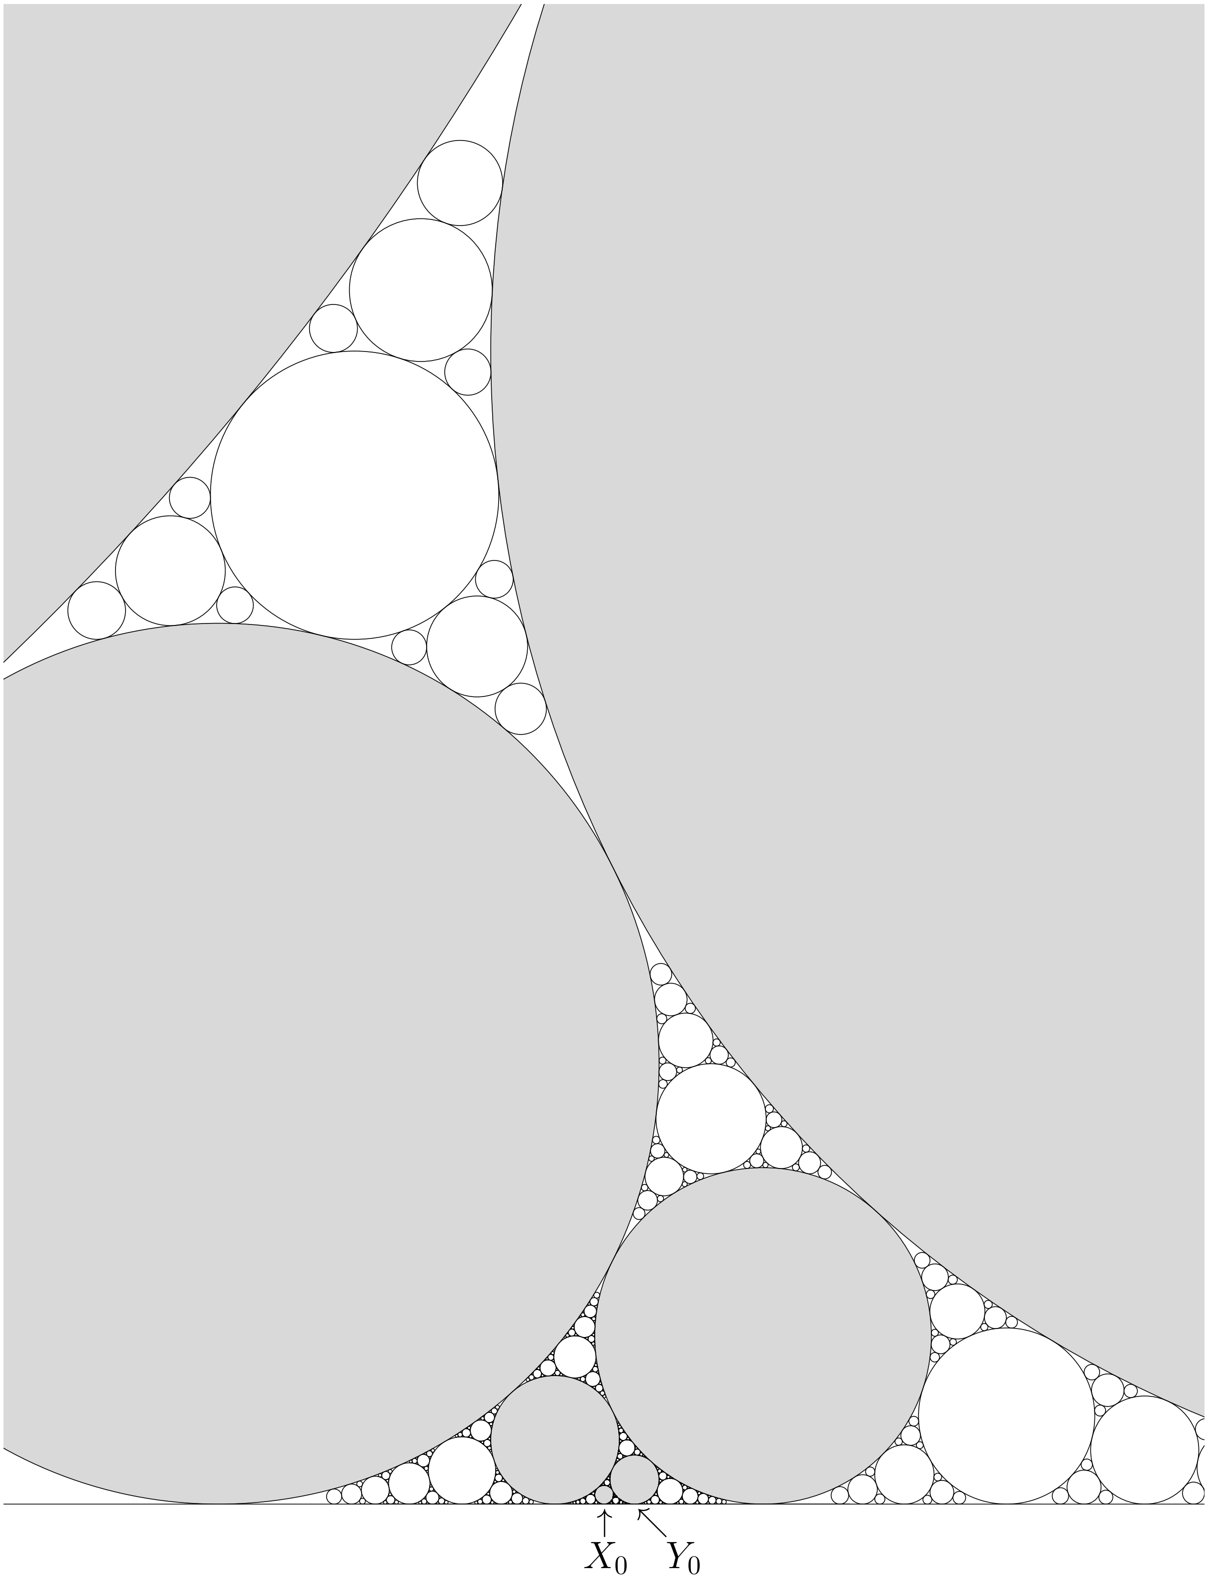

FIGURE 11. Eight generations of the packing $\mathcal{P}_\alpha$, where $\alpha = [\overline{1}] = \frac{1+\sqrt{5}}{2}$. The circles corresponding to the circle replacement algorithm are shaded gray.

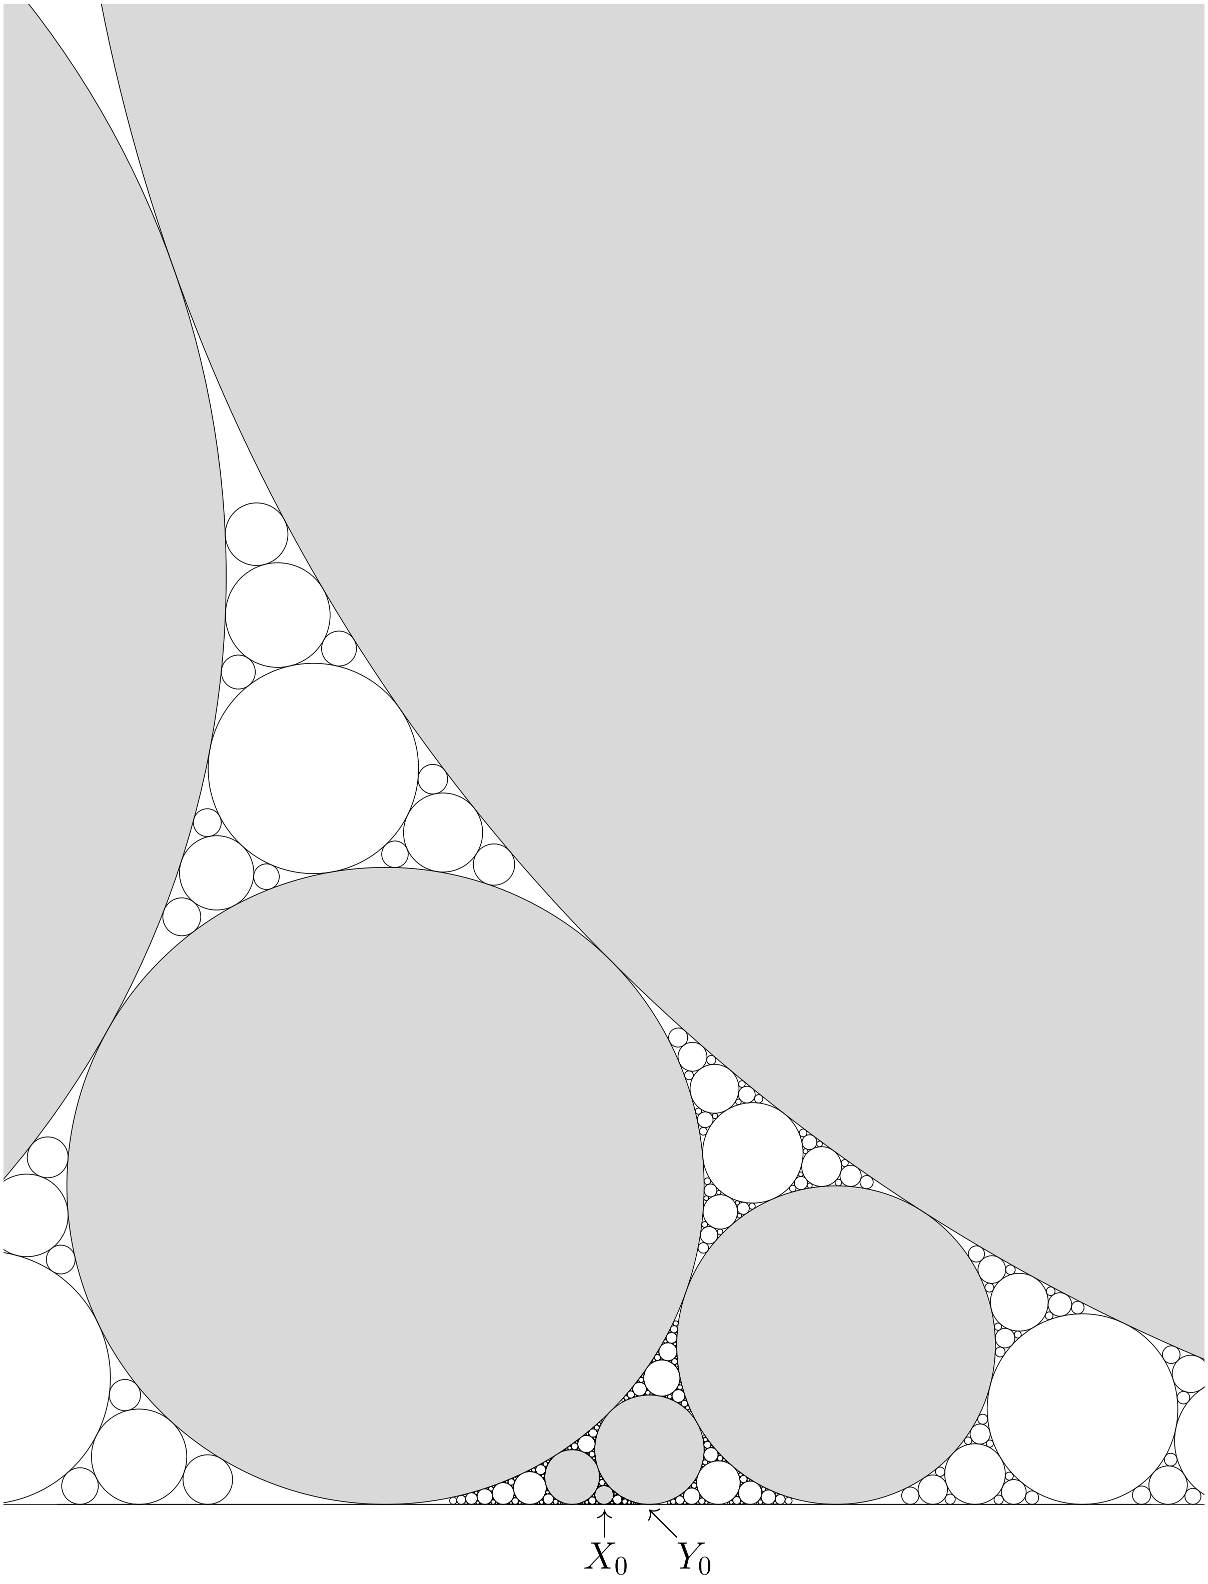

FIGURE 12. Eight generations of the packing $\mathcal{P}_\alpha$, where $\alpha = [\overline{2}] = 1 + \sqrt{2}$. The circles corresponding to the circle replacement algorithm are shaded gray.

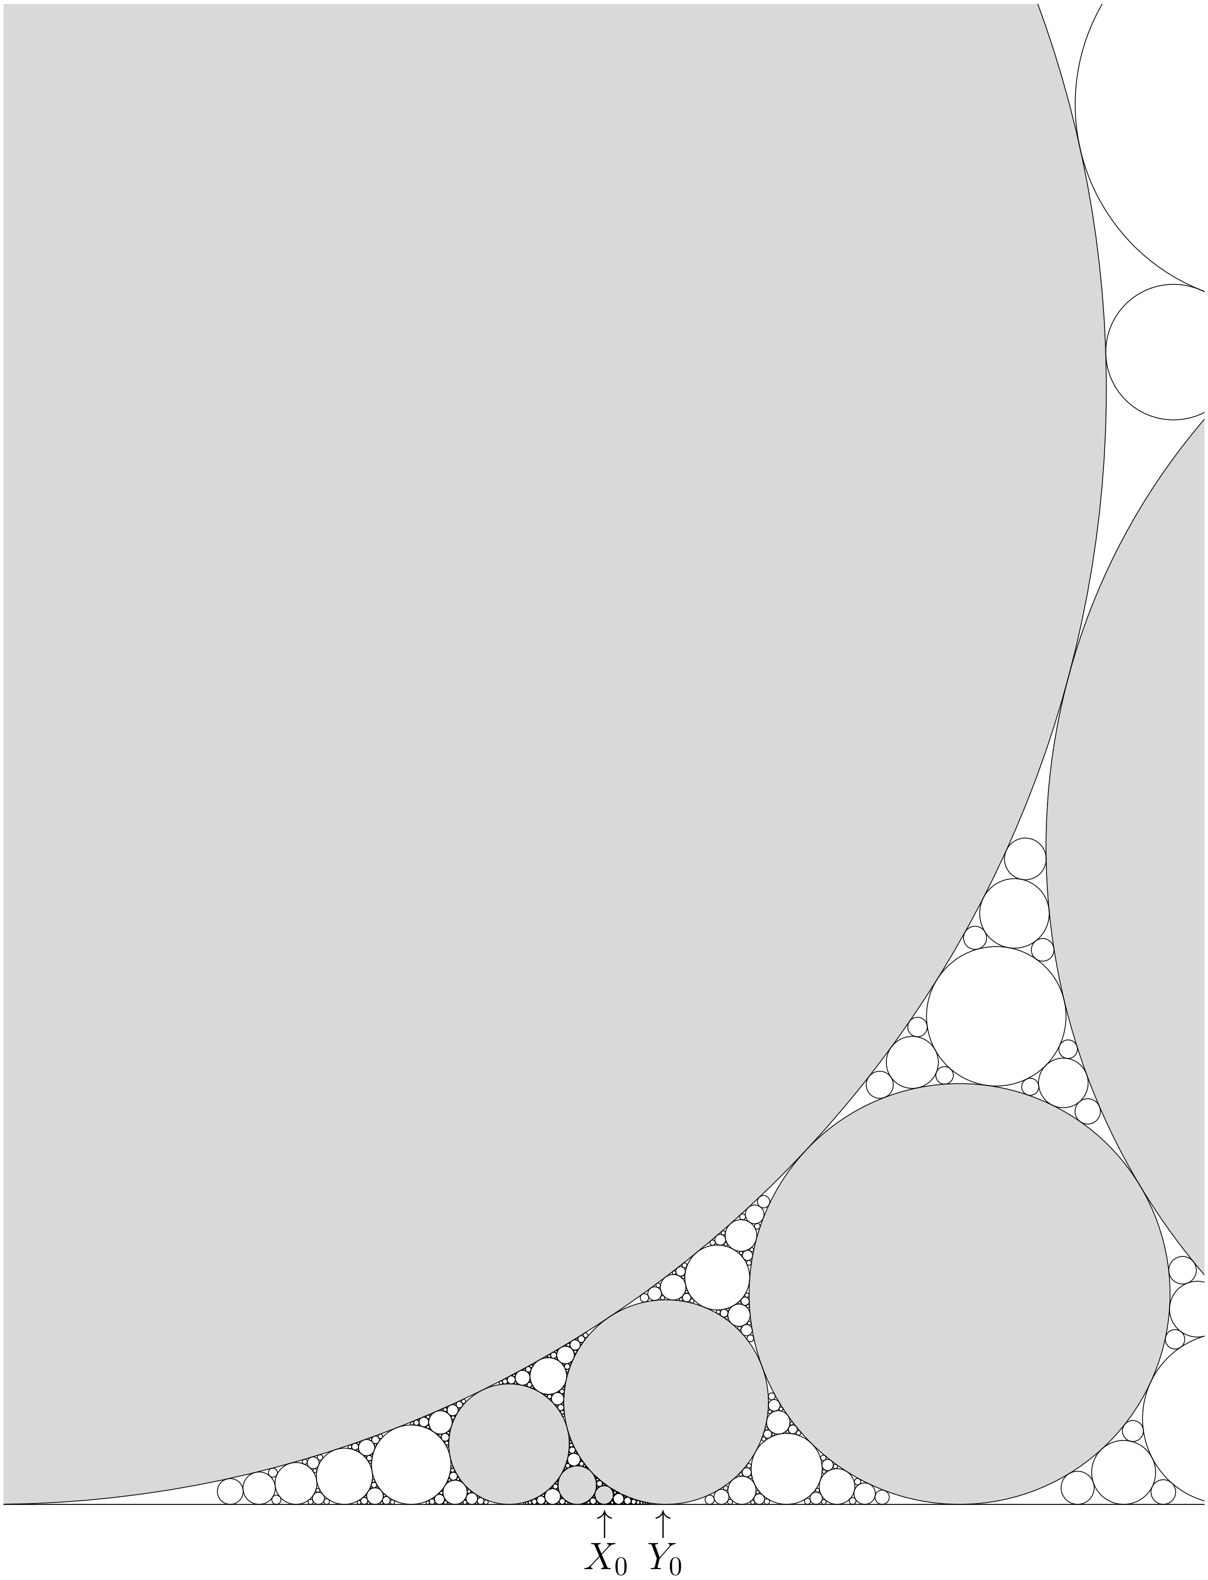

FIGURE 13. Eight generations of the packing $\mathcal{P}_\alpha$, where $\alpha = [\overline{3}] = \frac{3+\sqrt{13}}{2}$. The circles corresponding to the circle replacement algorithm are shaded gray.

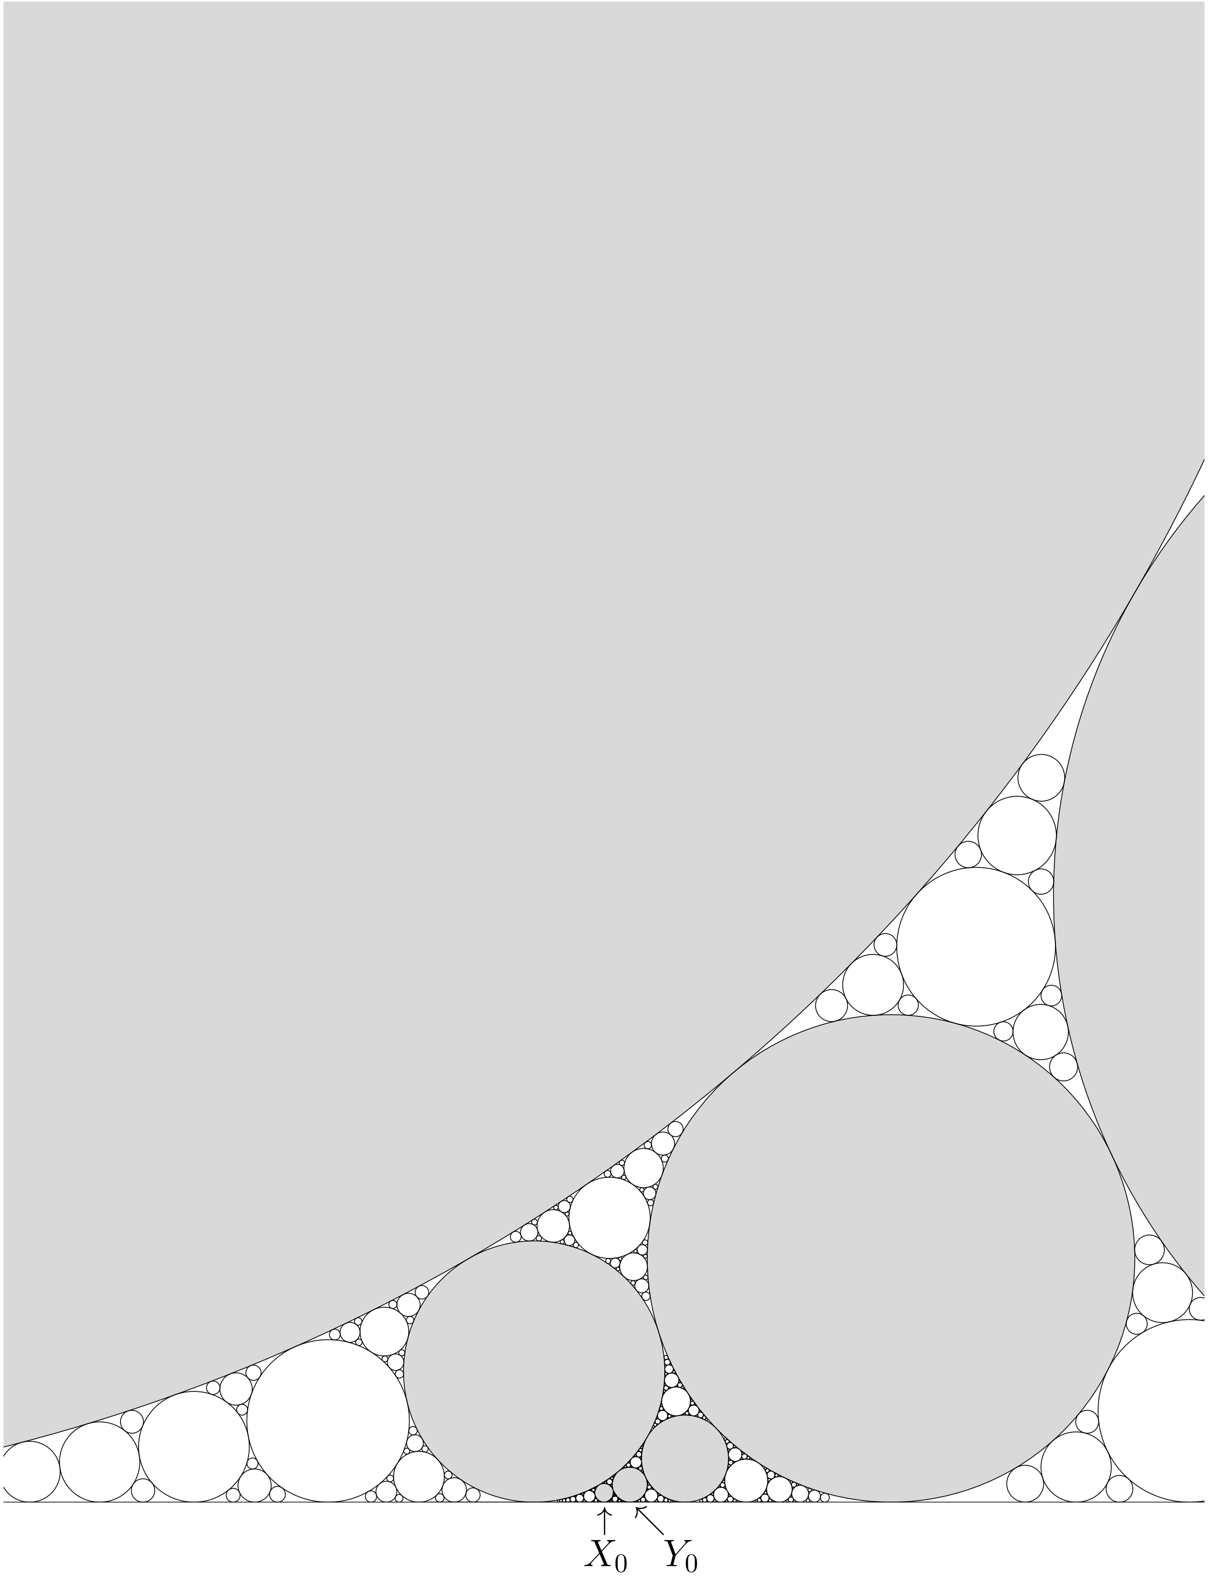

Figure 14. Eight generations of the packing $\mathcal{P}_\alpha$, where $\alpha = [\overline{1,2}] = \frac{1+\sqrt{3}}{2}$. The circles corresponding to the circle replacement algorithm are shaded gray.

11-14 are generated by those arising from the circle replacement algorithm. We can make this more precise as follows.

Suppose for simplicity that α has a *purely periodic* continued fraction expansion:

$$\alpha = [a_0, \dots, a_{n-1}, a_0, \dots, a_{n-1}, a_0, \dots] = [\overline{a_0, \dots, a_{n-1}}].$$

(By Theorem 4.7, we lose no generality in doing so.) After performing $N = a_0 + \cdots + a_{n-1} + n$ steps of the continued fraction algorithm, the initial segment consisting of the periodic part of the continued fraction expansion is removed; therefore $\alpha_N = \alpha$. It follows by Lemma 4.2 that applying N steps of the circle replacement algorithm to the pair $(X_0, Y_0) = (C_{(1,0)}, C_{(0,1)})$ yields a pair (X_N, Y_N) whose curvatures are in the same ratio (i.e., α^2) as the original circles (X_0, Y_0). There is therefore a similarity of $\mathcal{P}_\alpha$ that maps X_0 to X_N and Y_0 to Y_N by Lemma 3.5. (Note that if the continued fraction expansion of α is not purely periodic, the above argument determines instead a similarity mapping (X_M, Y_M) to (X_N, Y_N) for some M, N.) In fact, this argument may be repeated to show that there is a similarity of $\mathcal{P}_\alpha$ that maps X_0 to X_{kN} and Y_0 to Y_{kN} for each $k \in \mathbb{Z}_{\geq 0}$. The next lemma tells us that every self-similarity of $\mathcal{P}_\alpha$ arises from the continued fraction expansion of α in this way.

Lemma 4.9. *Suppose α is the positive real number with periodic continued fraction expansion $[\overline{a_0, \dots, \alpha_{n-1}}]$ and that this is the minimal periodic part. Let Φ denote the generator of* $\mathrm{Symm}(\mathcal{P}_\alpha) \cong \mathbb{Z}$ *for which Φ has scale factor greater than 1, and let $(X_0, Y_0) = (C_{(1,0)}, C_{(0,1)})$ be the generators for $\mathcal{P}_\alpha$. Then*

$$\Phi^k(X_0) = X_{kN}\ ,\ \Phi^k(Y_0) = Y_{kN}$$

for each $k \geq 0$, where $N = a_0 + \cdots + a_{n-1} + n$ as above, and Φ^k denotes the k-fold composition of Φ with itself (and Φ^0 is the identity map on $\mathbb{R}^2$).

Proof. We have already argued that for each $k \geq 0$ there is a self-similarity Φ_k of $\mathcal{P}_\alpha$ such that

$$\Phi_k(X_0) = X_{kN} \quad \text{and} \quad \Phi_k(Y_0) = Y_{kN}.$$

Furthermore, if N does not divide m, then $\alpha_m \neq \alpha$, so that the *only* self-similarities that correspond to the circle replacement algorithm are the Φ_k.

Because the radius of Y_0 is 1, the scale factor μ_k of Φ_k is equal to the radius of $\Phi_k(Y_0) = Y_{kN}$, which increases as k increases by the definition of the circle replacement algorithm. (Note that the Y_{kN} must necessarily be distinct.) Similarly, since the scale factor for Φ^k is μ^k, with $\mu > 1$, the scale factor for Φ^k also increases as k increases. It will therefore suffice to show that each self-similarity ϕ with scale factor greater than 1 satisfies

$$\phi(X_0) = X_m \quad \text{and} \quad \phi(Y_0) = Y_m$$

for some $m \geq 1$.

By Lemma 4.5, we may reduce this problem to showing that each of X_0 and Y_0 either lies in the bounded interstice for $\phi(X_0)$ and $\phi(Y_0)$ or is equal to one of $\phi(X_0)$ and $\phi(Y_0)$. Because self-similarities preserve the basic structure of the packing, X_0 (resp. Y_0) lies in the bounded interstice for $\phi(X_0)$ and $\phi(Y_0)$ if and only if $\phi^{-1}(X_0)$ (resp. $\phi^{-1}(Y_0)$) lies in the bounded interstice for X_0 and Y_0. It is therefore enough to show that if ψ is a self-similarity of $\mathcal{P}_\alpha$ with scale factor *less than* 1, then $\psi(X_0)$ lies in the bounded interstice for X_0 and Y_0. In that case $\psi(Y_0)$ will have to either be in the bounded interstice, or be equal to X_0 or Y_0.

By the proof of Theorem 3.22, for any nontrivial self-similarity ψ of $\mathcal{P}_\alpha$, we have

$$\psi(X_0) = \psi(C_{(1,0)}) = C_{\left(\frac{x-yq}{2}, -yr\right)}$$
$$\psi(Y_0) = \psi(C_{(0,1)}) = C_{\left(yp, \frac{x+yq}{2}\right)},$$

where x and y are integers satisfying $|x^2 - Dy^2| = 4$ with $y \neq 0$. Here $f(x) = px^2 + qx + r$, $p > 0$, is the primitive integer polynomial satisfied by α, and $D = q^2 - 4pr$ is the discriminant of f.

To show that $\psi(C_{(1,0)})$ either lies in the bounded interstice for $C_{(1,0)}$ and $C_{(0,1)}$, or is equal to $C_{(0,1)}$, it will suffice by Lemmas 2.15 and 2.20 to show that

$$\frac{x - yq}{2} > 0 \quad \text{and} \quad -yr > 0. \tag{4.10}$$

To see this, we first recall a result due to Galois [1, 2] concerning purely periodic continued fractions. Because α has a purely periodic continued fraction expansion, it is a reduced quadratic number; i.e., $\alpha > 1$ and $-1 < \alpha' < 0$, where α' is the quadratic conjugate of α. Since $q = -p(\alpha + \alpha')$ and $r = p\alpha\alpha'$, and since $p > 0$, it follows that $q, r < 0$. It is now sufficient to show that $x > 0$ and $y > 0$.

First note that we cannot have both $x, y < 0$ since then the label on $\psi(C_{(1,0)})$ would consist of two negative numbers, which is impossible by Proposition 2.21. We now show also that x and y cannot have different signs.

Let μ' be the scale factor of ψ which, by assumption, is less than 1. Since $\operatorname{curv}(C_{(0,1)}) = 1$ and $\operatorname{curv}(\psi(C_{(0,1)})) = (yp\alpha + \frac{x+yq}{2})^2$, it follows that

$$\mu' = \frac{1}{\left(yp\alpha + \frac{x+yq}{2}\right)^2}.$$

Therefore

$$\left| yp\alpha + \frac{x + yq}{2} \right| > 1.$$

Since α has purely periodic continued fraction expansion, it is greater than its conjugate, so we have $\alpha = \frac{-q+\sqrt{D}}{2p}$. The above inequality then implies

$$|x + y\sqrt{D}| > 2.$$

Now $x^2 - Dy^2 = \pm 4$, so $x = \pm\sqrt{Dy^2 \pm 4}$. If x and y have different signs, then this inequality becomes

$$|\sqrt{Dy^2 \pm 4} - \sqrt{Dy^2}| > 2.$$

But $|\sqrt{t+4} - \sqrt{t}| \leq 2$ for all $t \geq 0$, so in fact x and y must have the same sign (and $y \neq 0$ since we assumed that ψ was a nontrivial self-similarity). This completes the proof. □

Finally, we can also use continued fractions to see which packings have orientation-reversing self-similarities.

Theorem 4.11. *For a positive real number α, the circle packing $\mathcal{P}_\alpha$ has an orientation-reversing self-similarity if and only if the continued fraction expansion of α has odd period.*

Proof. We know from Lemma 4.9 that all self-similarities of $\mathcal{P}_\alpha$ correspond to periods in the continued fraction expansion of α. The self-similarity Φ is orientation-reversing if the corresponding pairs of circles in the circle replacement algorithm, say (X_n, Y_n) and (X_m, Y_m), for which $\Phi(X_n) = X_m$ and $\Phi(Y_n) = Y_m$, satisfy $X_n \prec Y_n$, but $Y_m \prec X_m$. Since the orientation of the circles X_k, Y_k changes once for each case of step (B) in the algorithm, that is, for each term in the continued fraction expansion, we see that Φ is orientation-reversing if and only if it corresponds to an odd period. □

In the examples, we see that the circle packings $\mathcal{P}_\alpha$ for $\alpha = [\bar{1}], [\bar{2}], [\bar{3}]$ do have orientation-reversing self-similarities, whereas that for $\alpha = [\overline{1, 2}]$ does not.

References

1. H. Davenport, *The Higher Arithmetic*, Cambridge University Press, 2003.
2. Évariste Galois, *Démonstration d'un théorème sur les fractions continues périodiques*, Ann. Math. Pures et Appl. **19** (1829), 294–301.
3. R. L. Graham, J. C. Lagarias, C. L. Mallows, A. R. Wilks, and C. H. Yan, *Apollonian Circle Packings: Number Theory*, J. Number Theory **100** (2003), 1–45.
4. ______, *Apollonian Circle Packings: Geometry and Group Theory I. The Apollonian Group*, Discrete Comput. Geom. **34** (2005), 547–585.
5. ______, *Apollonian Circle Packings: Geometry and Group Theory II. Super-Apollonian Group and Integral Packings*, Discrete Comput. Geom. **35** (2006), 1–36.
6. ______, *Apollonian Circle Packings: Geometry and Group Theory III. Higher Dimensions*, Discrete Comput. Geom. **35** (2006), 37–72.
7. G. H. Hardy and E. M. Wright, *An Introduction to the Theory of Numbers*, sixth ed., Oxford University Press, 2008.
8. Michael Jacobson and Hugh Williams, *Solving the Pell Equation*, CMS Books in Mathematics, Springer, 2009.
9. Daniel Pedoe, *On a Theorem in Geometry*, Amer. Math. Monthly **74** (1967), 627–640.